



(12) **United States Patent**
Hayden

(10) **Patent No.:** **US 9,087,024 B1**
(45) **Date of Patent:** **Jul. 21, 2015**

(54) **NARRATION OF NETWORK CONTENT**

(75) Inventor: **Andrew Hayden**, Seattle, WA (US)

(73) Assignee: **Amazon Technologies, Inc.**, Reno, NV (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 438 days.

6,430,624 B1 8/2002 Jamtgaard et al.
6,549,941 B1 4/2003 Jaquith et al.
6,560,620 B1 5/2003 Ching
6,625,624 B1 9/2003 Chen et al.
6,704,024 B2 3/2004 Robotham et al.
6,785,864 B1 8/2004 Te et al.
6,871,236 B2 3/2005 Fishman et al.

(Continued)

FOREIGN PATENT DOCUMENTS

WO WO 2013/003631 A2 1/2013

(21) Appl. No.: **13/359,028**

(22) Filed: **Jan. 26, 2012**

(51) **Int. Cl.**

G06F 17/22 (2006.01)

G06F 17/21 (2006.01)

G06F 3/0481 (2013.01)

G10L 25/30 (2013.01)

G10L 15/30 (2013.01)

(52) **U.S. Cl.**

CPC **G06F 17/211** (2013.01); **G06F 3/0481** (2013.01); **G10L 15/30** (2013.01); **G10L 25/30** (2013.01)

(58) **Field of Classification Search**

USPC 715/200, 201, 202, 205, 234, 249, 255,
715/256, 273, 704, 723, 727, 746, 760;
704/2, 214, 254, 269, 5, E15.001

See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,634,064 A 5/1997 Warnock et al.
5,872,850 A 2/1999 Klein et al.
5,930,752 A * 7/1999 Kawaguchi et al. 704/270.1
5,961,593 A 10/1999 Gabber et al.
6,049,812 A 4/2000 Bertram et al.
6,084,582 A * 7/2000 Qureshi et al. 715/243
6,108,637 A 8/2000 Blumenau
6,138,156 A 10/2000 Fletcher et al.
6,195,679 B1 2/2001 Bauersfeld et al.

OTHER PUBLICATIONS

Baumann, A., et al., Enhancing STEM Classes Using Weave: A Collaborative Web-Based Visualization Environment, Integrated Stem Education Conference, Apr. 2, 2011, Ewing, New Jersey, pp. 2A-1-2A-4.

(Continued)

Primary Examiner — Maikhanh Nguyen

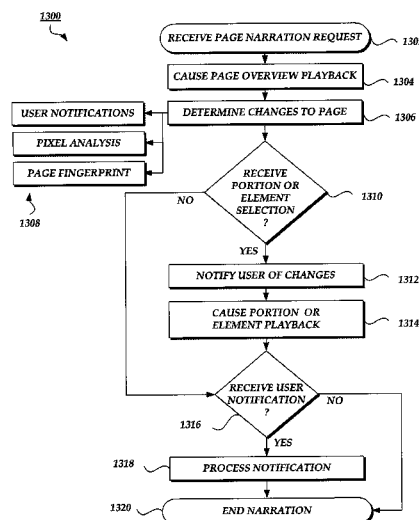
(74) *Attorney, Agent, or Firm* — Knobbe, Martens, Olson & Bear

(57)

ABSTRACT

A process is directed to the generation, maintenance, and provision of narrations for requested network content. A client computing device transmits a narration of a network resource, such as a web page, to a network computing provider. The narration can be separated into portions, each portion corresponding to a portion of the network resource. A second client computing device can request the network resource and the narration. The network computing provider can determine whether there has been any change in the requested network resource since the narration was received, and can modify the narration provided to the second client computing device accordingly. The second client computing device can playback appropriate portions of the narration while still being capable of regular interaction with the requested network resource.

24 Claims, 16 Drawing Sheets



(56)

References Cited

U.S. PATENT DOCUMENTS

- 6,944,665 B2 9/2005 Brown et al.
 6,963,850 B1 11/2005 Bezos et al.
 7,003,442 B1 2/2006 Tsuda
 7,051,084 B1 5/2006 Hayton et al.
 7,054,952 B1 5/2006 Schwerdtfeger et al.
 7,085,736 B2 8/2006 Keezer et al.
 7,159,023 B2 1/2007 Tufts
 7,171,478 B2 1/2007 Lueckhoff et al.
 7,191,211 B2 3/2007 Tuli
 7,219,136 B1* 5/2007 Danner et al. 709/219
 7,274,775 B1* 9/2007 Zavaliagos et al. 379/75
 7,353,252 B1 4/2008 Yang et al.
 7,373,313 B1 5/2008 Kahle et al.
 7,539,619 B1* 5/2009 Seligman et al. 704/277
 7,543,059 B2 6/2009 Johnson et al.
 7,653,544 B2* 1/2010 Bradley et al. 704/260
 7,792,944 B2 9/2010 DeSantis et al.
 7,831,582 B1 11/2010 Scofield et al.
 7,966,184 B2* 6/2011 O'Connor et al. 704/260
 7,975,000 B2 7/2011 Dixon et al.
 7,996,912 B2 8/2011 Spalink et al.
 8,010,545 B2 8/2011 Stefik et al.
 8,015,496 B1 9/2011 Rogers
 8,060,463 B1 11/2011 Spiegel
 8,073,850 B1 12/2011 Hubbard et al.
 8,078,464 B2* 12/2011 Warford et al. 704/246
 8,099,284 B2* 1/2012 Arizmendi et al. 704/270.1
 8,103,742 B1 1/2012 Green
 8,185,621 B2 5/2012 Kasha
 8,204,359 B2* 6/2012 Ljolje 386/241
 8,249,904 B1 8/2012 DeSantis et al.
 8,271,887 B2 9/2012 Offer et al.
 8,307,396 B2* 11/2012 Tofighbakhsh 725/58
 8,316,124 B1 11/2012 Baumbach et al.
 8,336,049 B2 12/2012 Medovich
 8,751,023 B2* 6/2014 Lu et al. 700/94
 2001/0039490 A1 11/2001 Verbitsky et al.
 2002/0030703 A1 3/2002 Robertson et al.
 2002/0194302 A1 12/2002 Blumberg
 2003/0023712 A1 1/2003 Zhao et al.
 2003/0041106 A1 2/2003 Tuli
 2004/0006481 A1* 1/2004 Kieczka et al. 704/276
 2004/0083294 A1 4/2004 Lewis
 2004/0139208 A1 7/2004 Tuli
 2004/0181613 A1 9/2004 Hashimoto et al.
 2004/0205448 A1 10/2004 Grefenstette et al.
 2004/0220905 A1 11/2004 Chen et al.
 2004/0243622 A1 12/2004 Morisawa
 2005/0010863 A1 1/2005 Zernik
 2005/0033577 A1* 2/2005 Bradley et al. 704/270.1
 2005/0060643 A1 3/2005 Glass et al.
 2005/0138382 A1 6/2005 Hougaard et al.
 2005/0183039 A1 8/2005 Revis
 2005/0246193 A1 11/2005 Roever et al.
 2006/0085766 A1 4/2006 Dominowska et al.
 2006/0095336 A1 5/2006 Heckerman et al.
 2006/0122889 A1 6/2006 Burdick et al.
 2006/0168510 A1 7/2006 Bryar et al.
 2006/0184421 A1 8/2006 Lipsky et al.
 2006/0248442 A1 11/2006 Rosenstein et al.
 2006/0277167 A1 12/2006 Gross et al.
 2006/0294461 A1 12/2006 Nadamoto et al.
 2007/0022072 A1 1/2007 Kao et al.
 2007/0027672 A1 2/2007 Decary et al.
 2007/0094241 A1 4/2007 Blackwell et al.
 2007/0118801 A1* 5/2007 Harshbarger et al. 715/730
 2007/0124693 A1 5/2007 Dominowska et al.
 2007/0139430 A1 6/2007 Korn et al.
 2007/0208687 A1* 9/2007 O'Connor et al. 707/1
 2007/0240160 A1 10/2007 Paterson-Jones et al.
 2007/0288589 A1 12/2007 Chen et al.
 2008/0028334 A1 1/2008 De Mes
 2008/0104502 A1 5/2008 Olston
 2008/0183672 A1 7/2008 Canon et al.
 2008/0184128 A1 7/2008 Swenson et al.
 2008/0232775 A1* 9/2008 Ljolje 386/96
 2008/0320225 A1 12/2008 Panzer et al.
 2009/0012969 A1 1/2009 Rail et al.
 2009/0164924 A1 6/2009 Flake et al.
 2009/0204478 A1 8/2009 Kaib et al.
 2009/0217199 A1 8/2009 Hara et al.
 2009/0248680 A1 10/2009 Kalavade
 2009/0254867 A1 10/2009 Farouki et al.
 2009/0282021 A1 11/2009 Bennett
 2009/0287698 A1 11/2009 Marmaros et al.
 2009/0327914 A1 12/2009 Adar et al.
 2010/0036740 A1 2/2010 Barashi
 2010/0057639 A1 3/2010 Schwarz et al.
 2010/0125507 A1 5/2010 Tarantino, III et al.
 2010/0131594 A1 5/2010 Kashimoto
 2010/0138293 A1 6/2010 Ramer et al.
 2010/0218106 A1 8/2010 Chen et al.
 2010/0293190 A1 11/2010 Kaiser et al.
 2010/0305726 A1* 12/2010 Lu et al. 700/94
 2010/0312788 A1 12/2010 Bailey
 2010/0318892 A1 12/2010 Teevan et al.
 2010/0322042 A1* 12/2010 Serletic et al. 369/1
 2010/0332513 A1 12/2010 Azar et al.
 2011/0022957 A1 1/2011 Lee
 2011/0029854 A1 2/2011 Nashi et al.
 2011/0055203 A1 3/2011 Gutt et al.
 2011/0078140 A1 3/2011 Dube et al.
 2011/0078705 A1 3/2011 Maclinovsky et al.
 2011/0119661 A1 5/2011 Agrawal et al.
 2011/0161849 A1 6/2011 Stallings et al.
 2011/0173177 A1 7/2011 Junqueira et al.
 2011/0173637 A1 7/2011 Brandwine et al.
 2011/0178868 A1 7/2011 Garg et al.
 2011/0185025 A1 7/2011 Cherukuri et al.
 2011/0191327 A1 8/2011 Lee
 2011/0197121 A1 8/2011 Kletter
 2011/0246873 A1 10/2011 Tolle et al.
 2011/0289074 A1 11/2011 Leban
 2011/0296341 A1 12/2011 Koppert
 2011/0302510 A1 12/2011 Harrison et al.
 2011/0321100 A1* 12/2011 Tofighbakhsh 725/58
 2012/0072821 A1 3/2012 Bowling
 2012/0084644 A1 4/2012 Robert et al.
 2012/0096365 A1 4/2012 Wilkinson et al.
 2012/0110017 A1 5/2012 Gu et al.
 2012/0137201 A1 5/2012 White et al.
 2012/0143944 A1 6/2012 Reeves et al.
 2012/0150844 A1 6/2012 Lindahl et al.
 2012/0166922 A1 6/2012 Rolles
 2012/0198516 A1 8/2012 Lim
 2012/0215834 A1 8/2012 Chen et al.
 2012/0215919 A1 8/2012 Labat et al.
 2012/0284629 A1 11/2012 Peters et al.
 2012/0317295 A1 12/2012 Baird et al.
 2012/0331406 A1 12/2012 Baird et al.
 2013/0007101 A1 1/2013 Trahan et al.
 2013/0007102 A1 1/2013 Trahan et al.
 2013/0031461 A1 1/2013 Hou et al.
 2013/0080611 A1 3/2013 Li et al.

OTHER PUBLICATIONS

De Carvalho, L.G., et al., Synchronizing Web Browsing Data With Browser, Proceedings of the IEEE Symposium on Computers and Communications, Jun. 22-25, 2010, Riccione, Italy, pp. 738-743.
 Close 'n' Forget Firefox add on, Evilfantasy's blog, <http://evilfantasy.wordpress.com/2009/03/24/close-%E2%80%98n%E2%80%99-forget-firefox-add-on/>, retrieved Mar. 24, 2009, 1 page.
 Chen, H., et al., "Bringing Order to the Web: Automatically Categorizing Search Results," Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, Apr. 1-6, 2000, pp. 145-152.
 Rao, H.C.-H., et al., "A Proxy-Based Personal Web Archiving Service," Operating Systems Review, 35(1):61-72, 2001.
 Teevan, J., et al., "Changing How People View Changes on the Web," 2009, Proceedings of the 22nd Annual ACM Symposium on User Interface Software and Technology, New York, 2009, pp. 237-246.
 Bango, Rey "How JS & Ajax work in Opera Mini 4", Nov. 2, 2007, XP055050107, Retrieved from the Internet.

(56)

References Cited

OTHER PUBLICATIONS

Brinkmann, M, "Record and Share your browser history with Hooeey," ghacks.net, Feb. 26, 2008, 6 pages, printed on Jan. 25, 2013.
 Considine, A, "The Footprints of Web Feet," The New York Times, Mar. 4, 2011, 3 pages, printed on Jan. 25, 2013.
 EyeBrowse: Record, Visualize and Share your Browser History, Information Aesthetics, Sep. 18, 2009, 2 pages, printed on Jan. 25, 2013.
 Feuerstein, Adam, "Flyswat Takes Aim," San Francisco Business Times, printed from <http://www.bizjournals.com/sanfrancisco/stories/1999/10/25/story2.html?t=printable>, Oct. 22, 1999, 2 pages.
 Gabber et al., "How to Make Personalized Web Browsing Simple, Secure, and Anonymous," Financial Cryptography, 16 pages (1997).
 Gingerich, Jason, "Keycorp Making Site Into Portal," KRTBN

Knight-Ridder Tribune Business News (South Bend Tribune, Indiana), Oct. 25, 1999, 2 pages.
 Hopper, D. Ian, "Desktops Now Have Power to Comparison-Shop," Oct. 18, 1999, printed from <http://www.cnn.com/TECH/computing/9910/18/r.u.sure/index.html>, 3 pages.
 Van Kleek, M, Introducing "Eyebrowse"—Track and share your web browsing in real time, Haystack Blog, Aug. 28, 2009, 3 pages, printed on Jan. 25, 2013.
 Web page titled "RSS Ticker: Add-ons for Firefox," <https://addons.mozilla.org/en-US/firefox/addon/rss-ticker/>, 3 printed pages, printed on Feb. 7, 2013.
 Web page titled "What Internet Users Do On A Typical Day, Trend Data (Adults), Pew Internet & American Life Project," printed from [http://pewinternet.org/Static-Pages/Trend-Data-\(Adults\)/Online-Activities-Daily.aspx](http://pewinternet.org/Static-Pages/Trend-Data-(Adults)/Online-Activities-Daily.aspx) on Nov. 29, 2012, 4 pages.

* cited by examiner

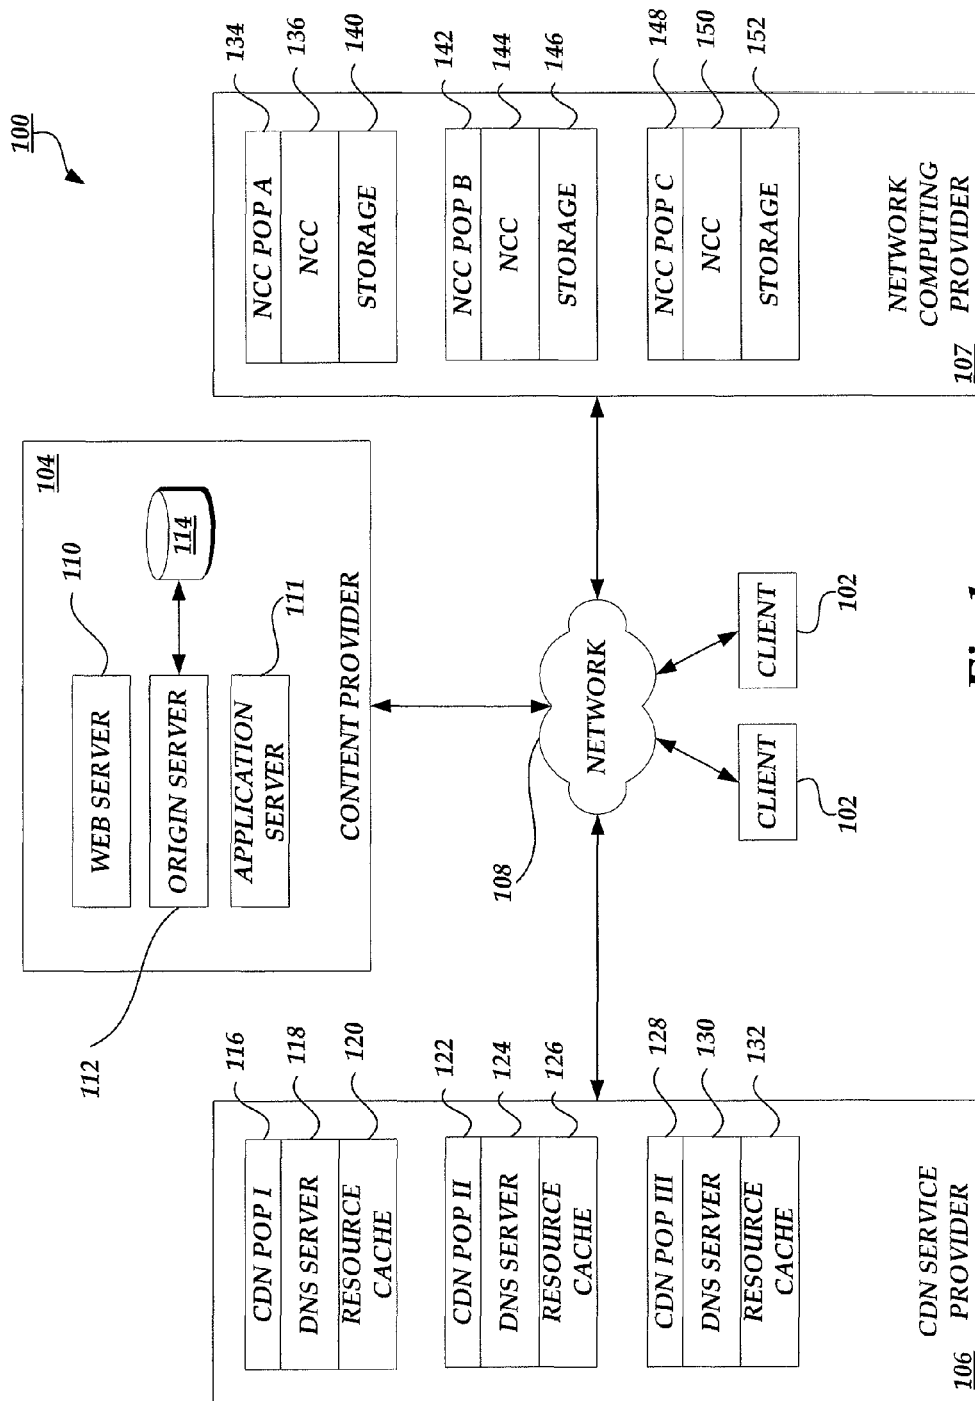


Fig. 1.

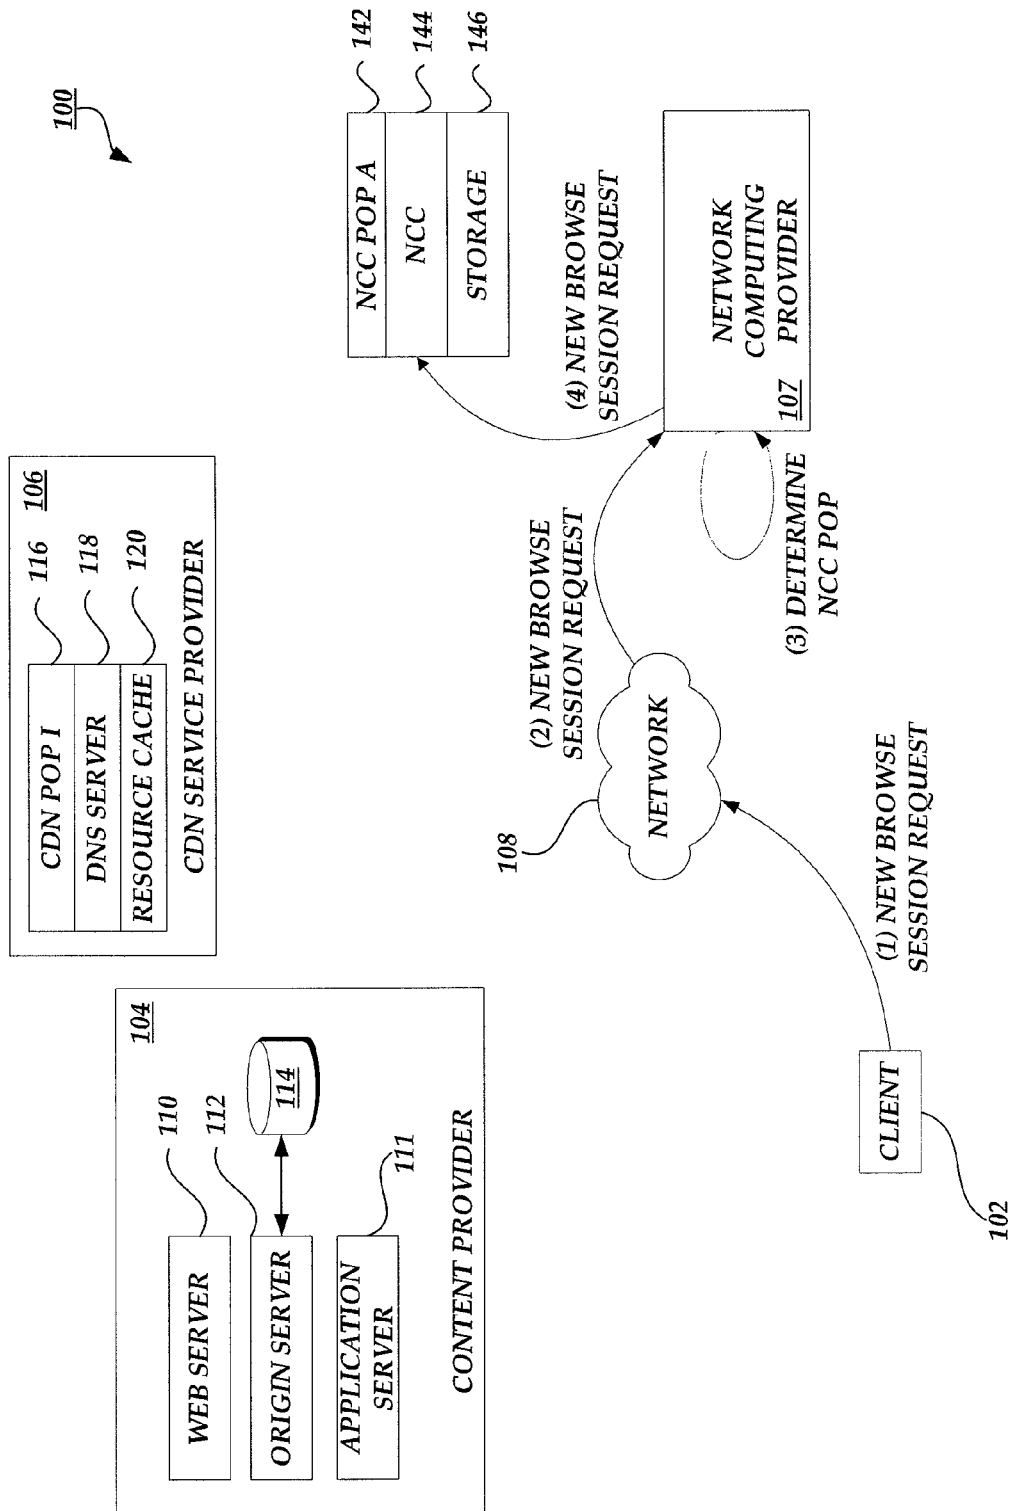


Fig. 2.

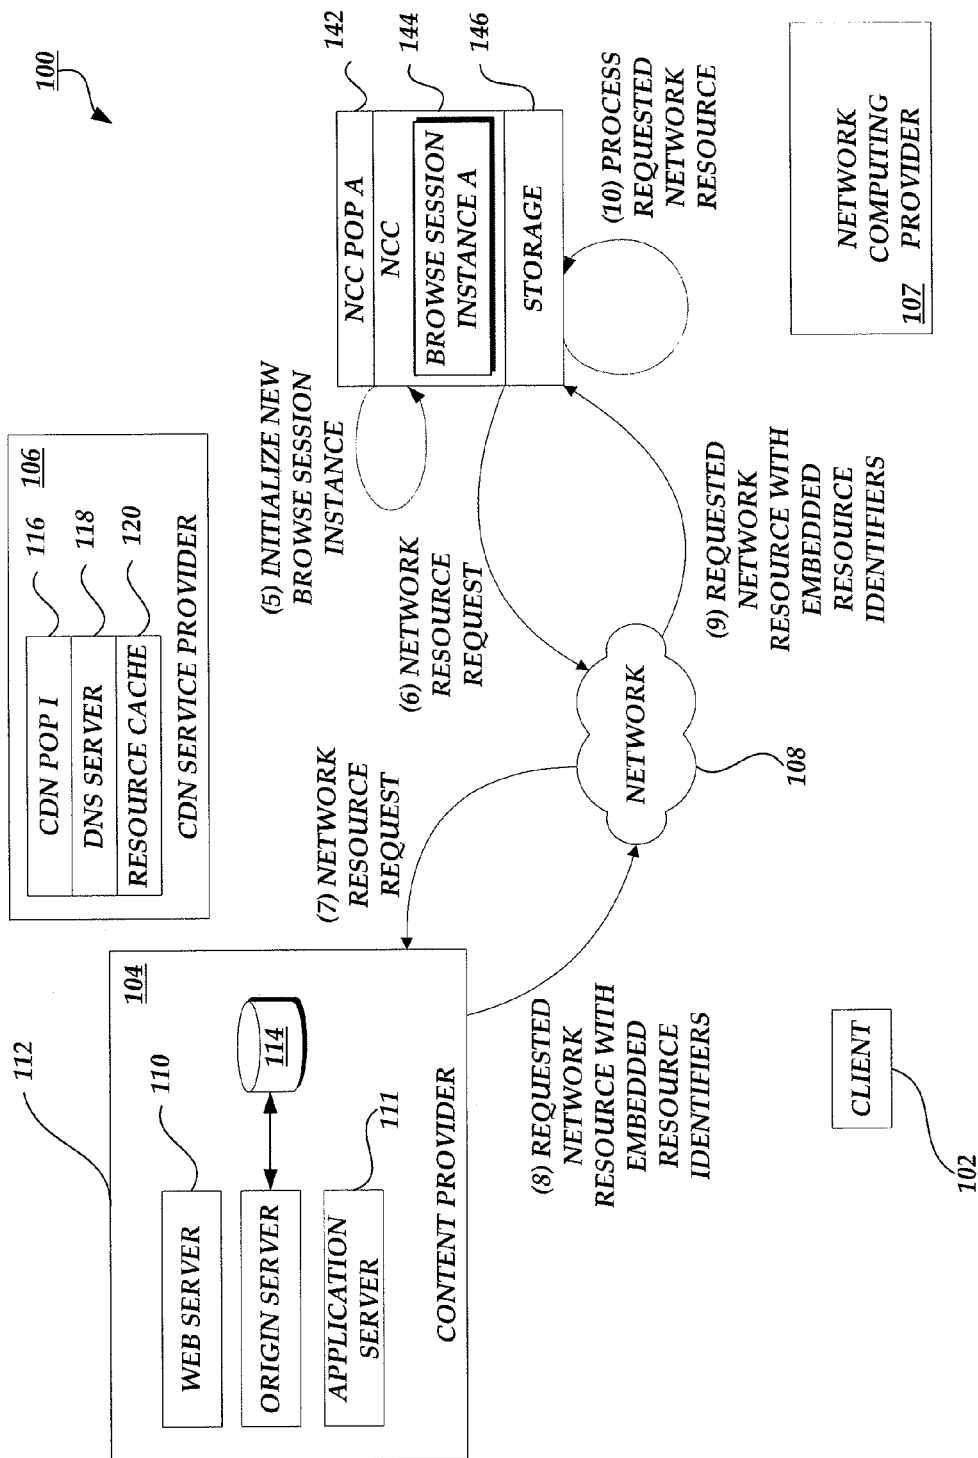


Fig. 3.

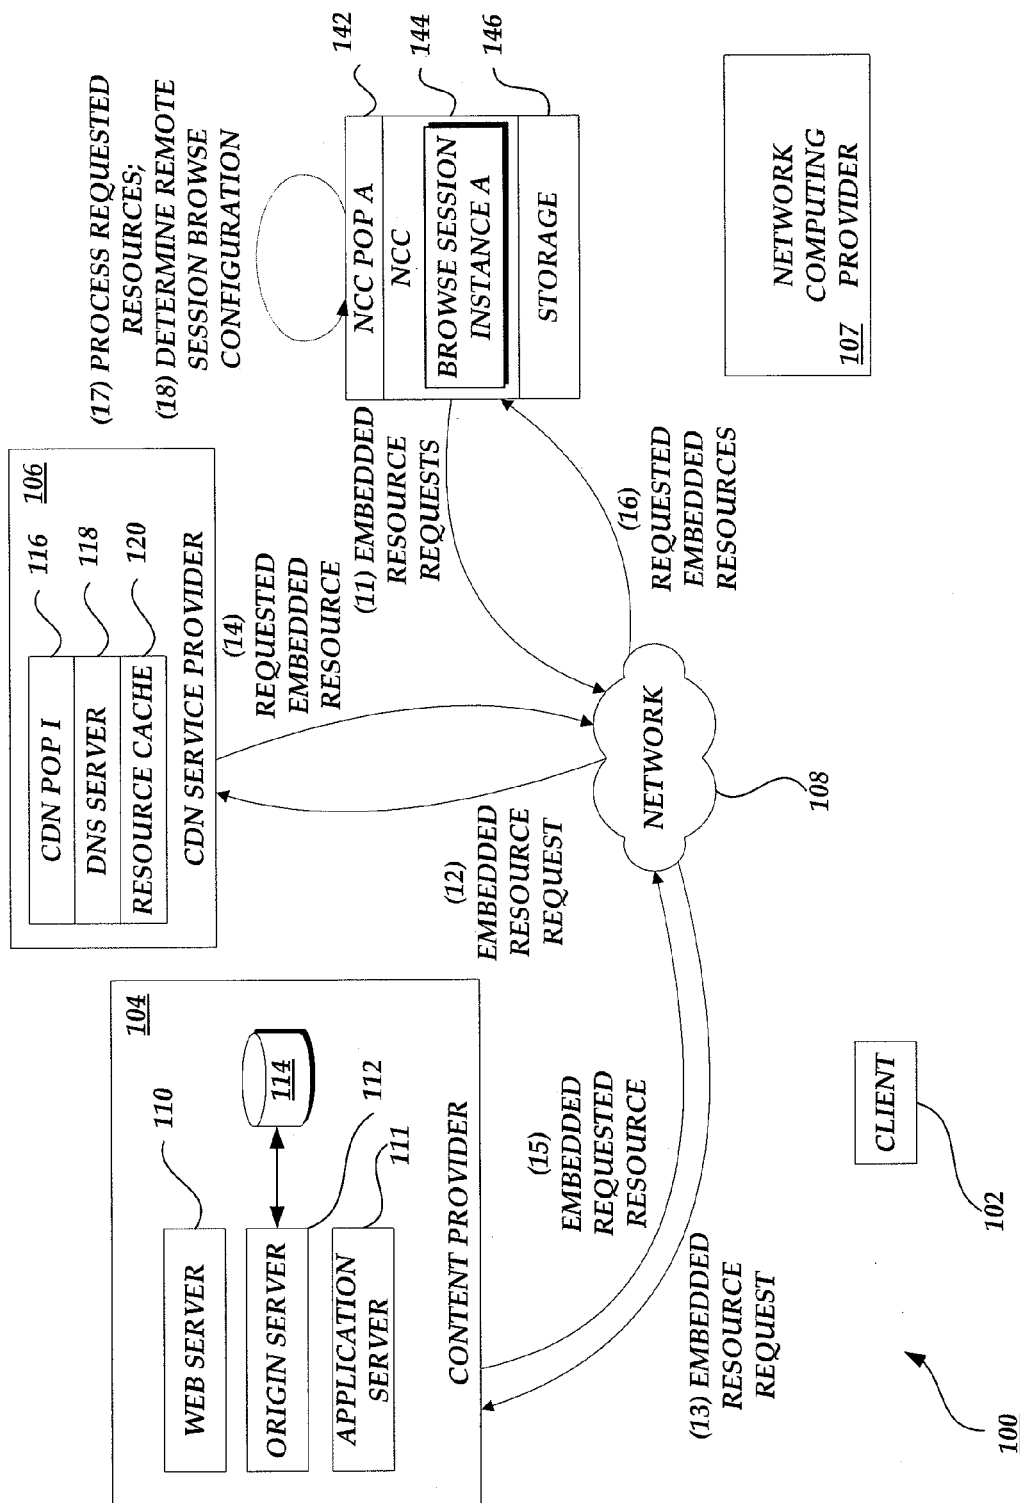


Fig. 4.

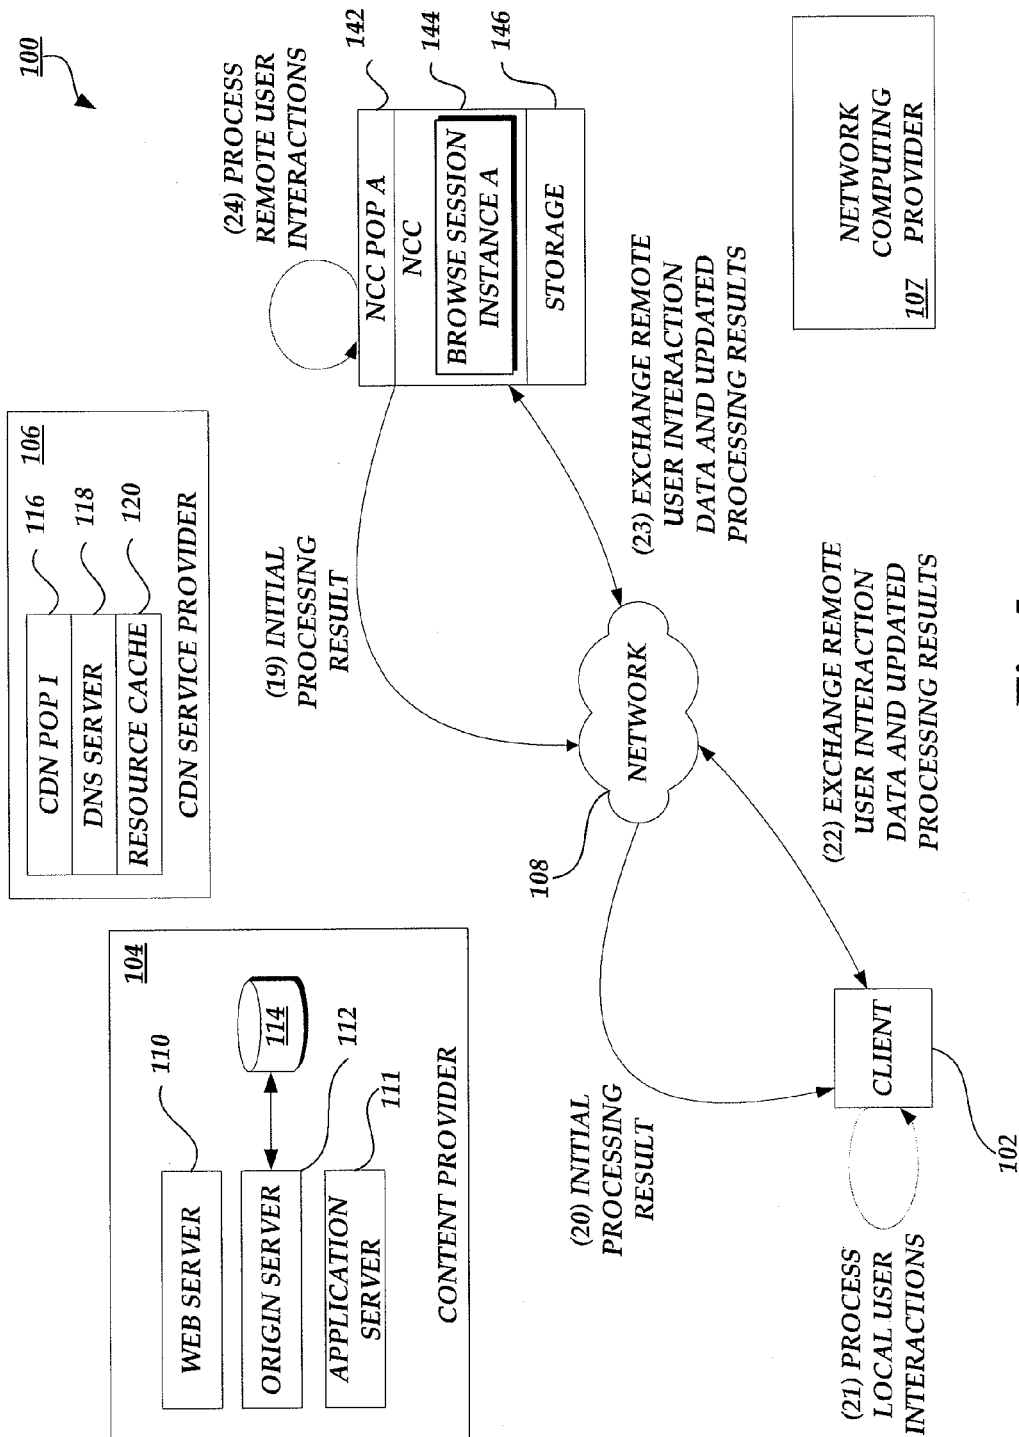


Fig. 5.

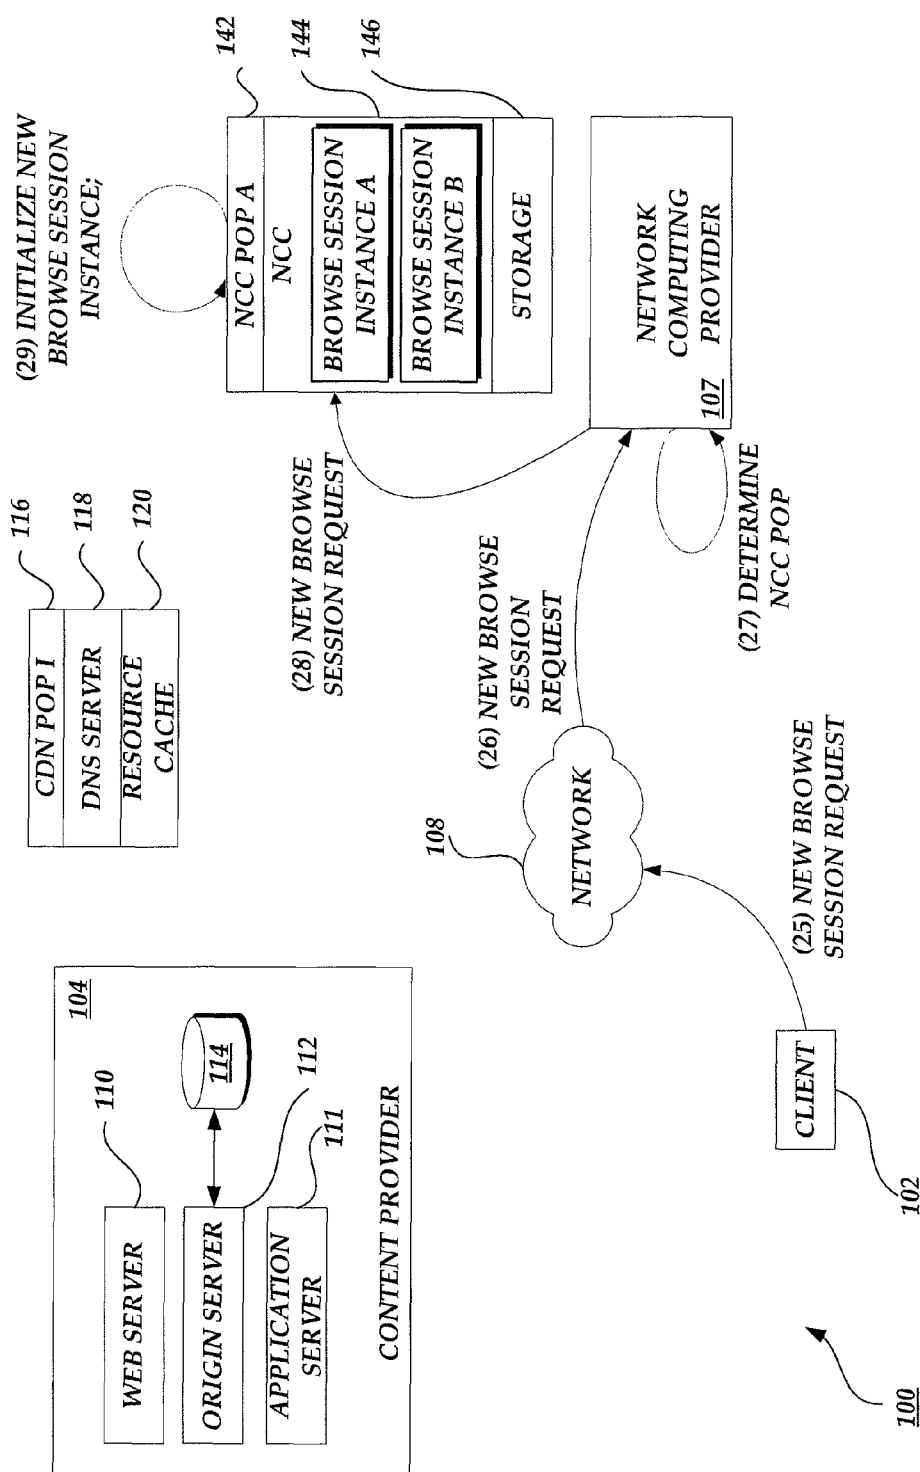
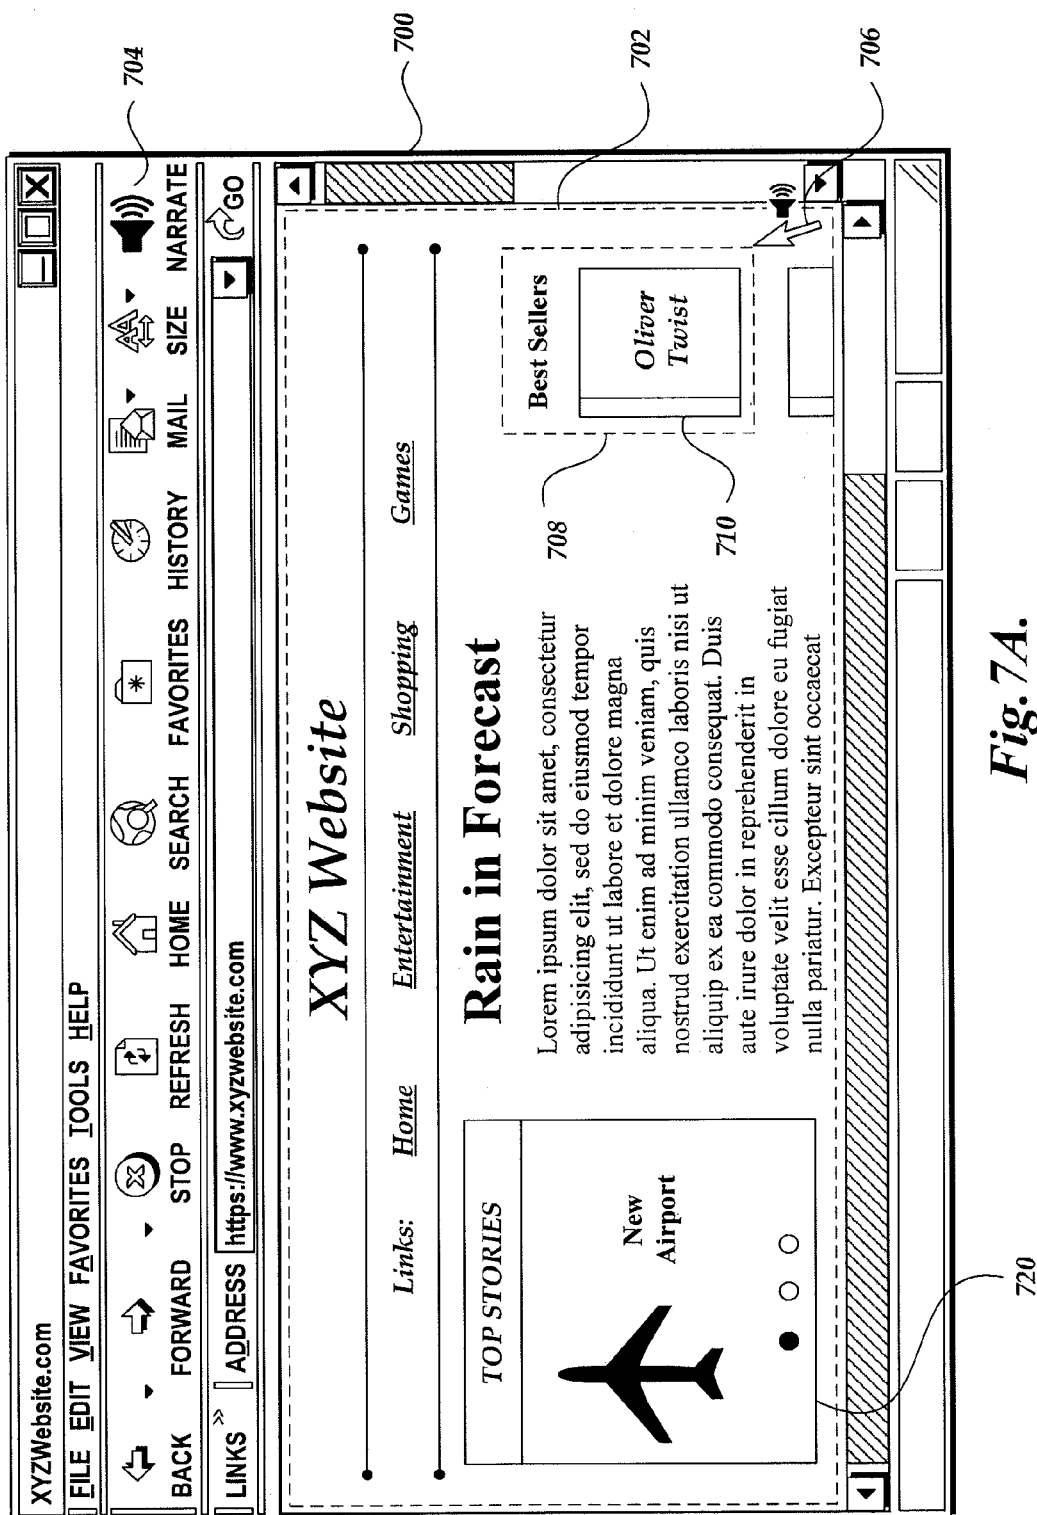


Fig. 6.



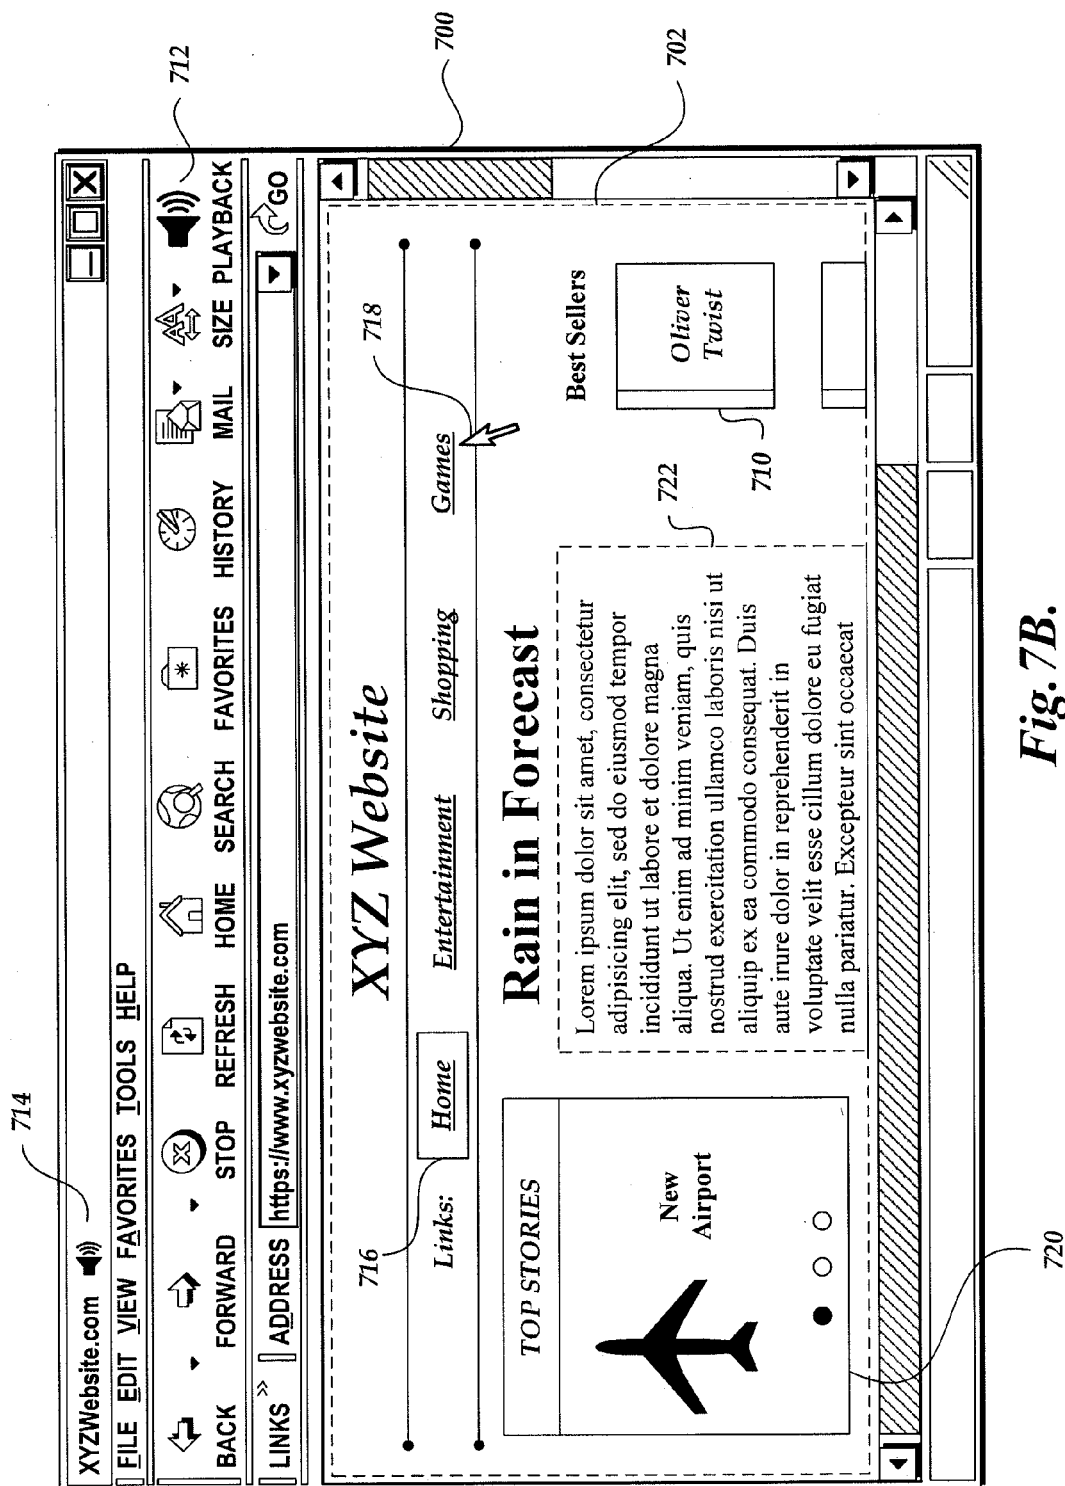
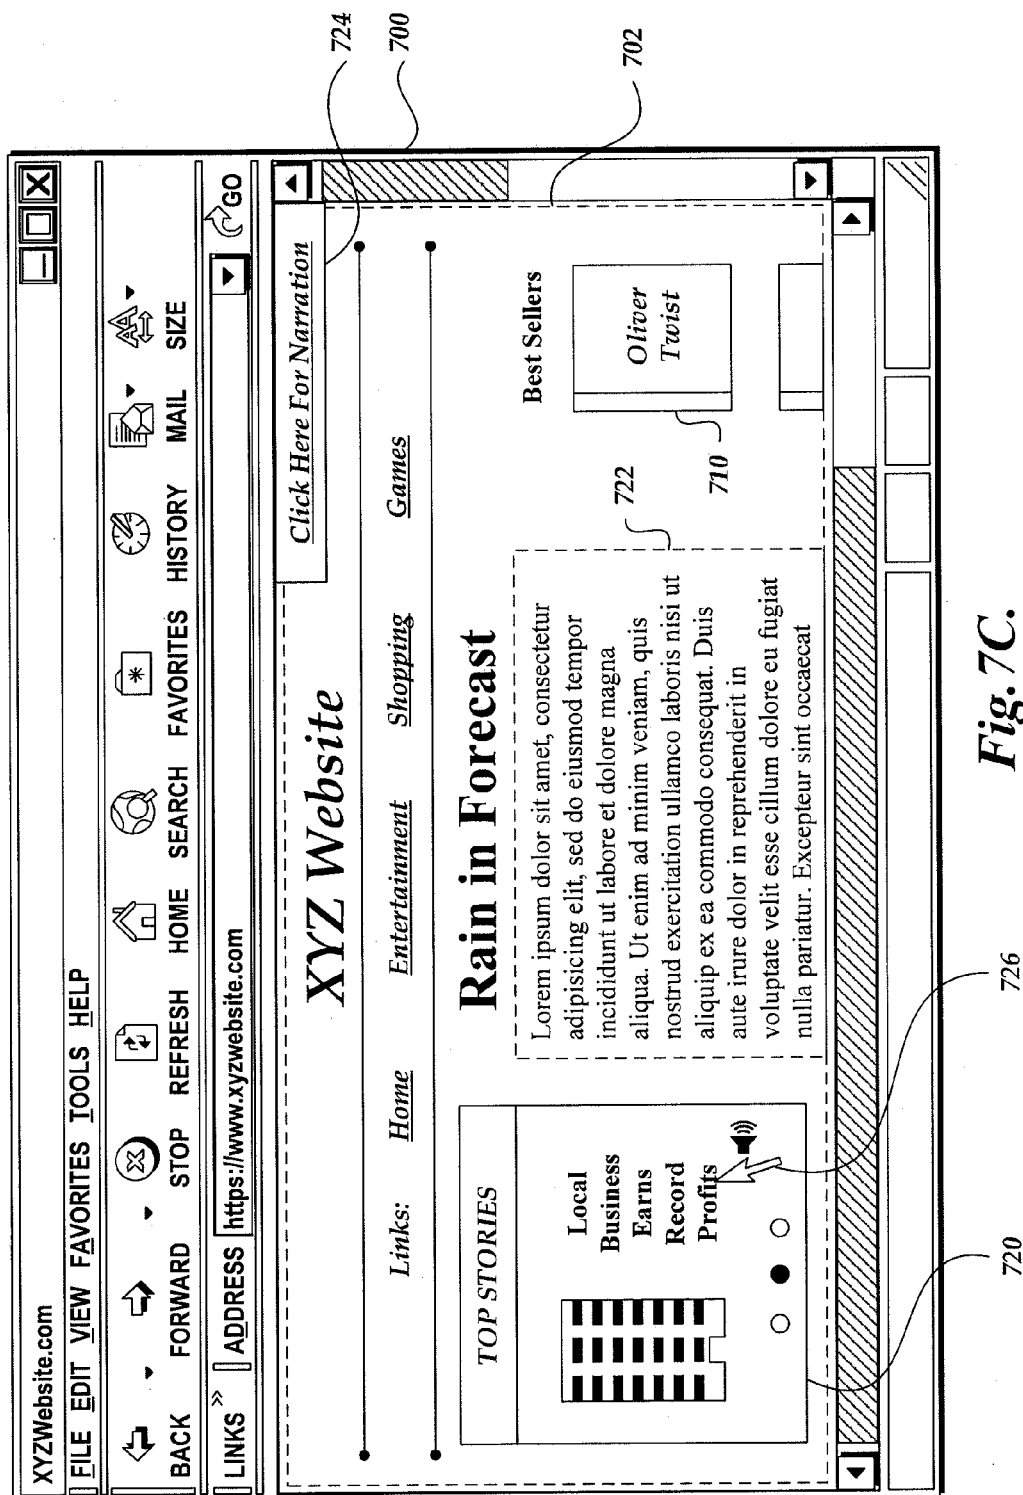
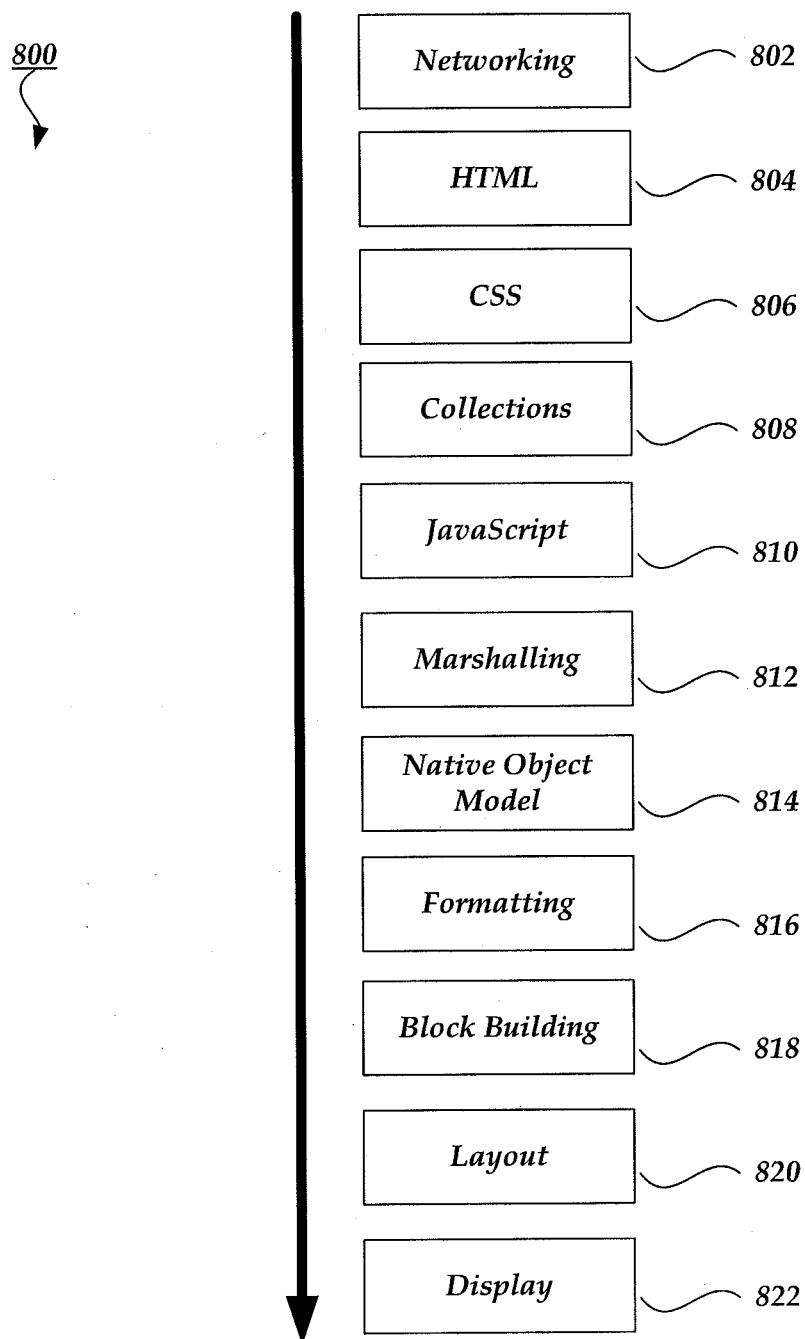
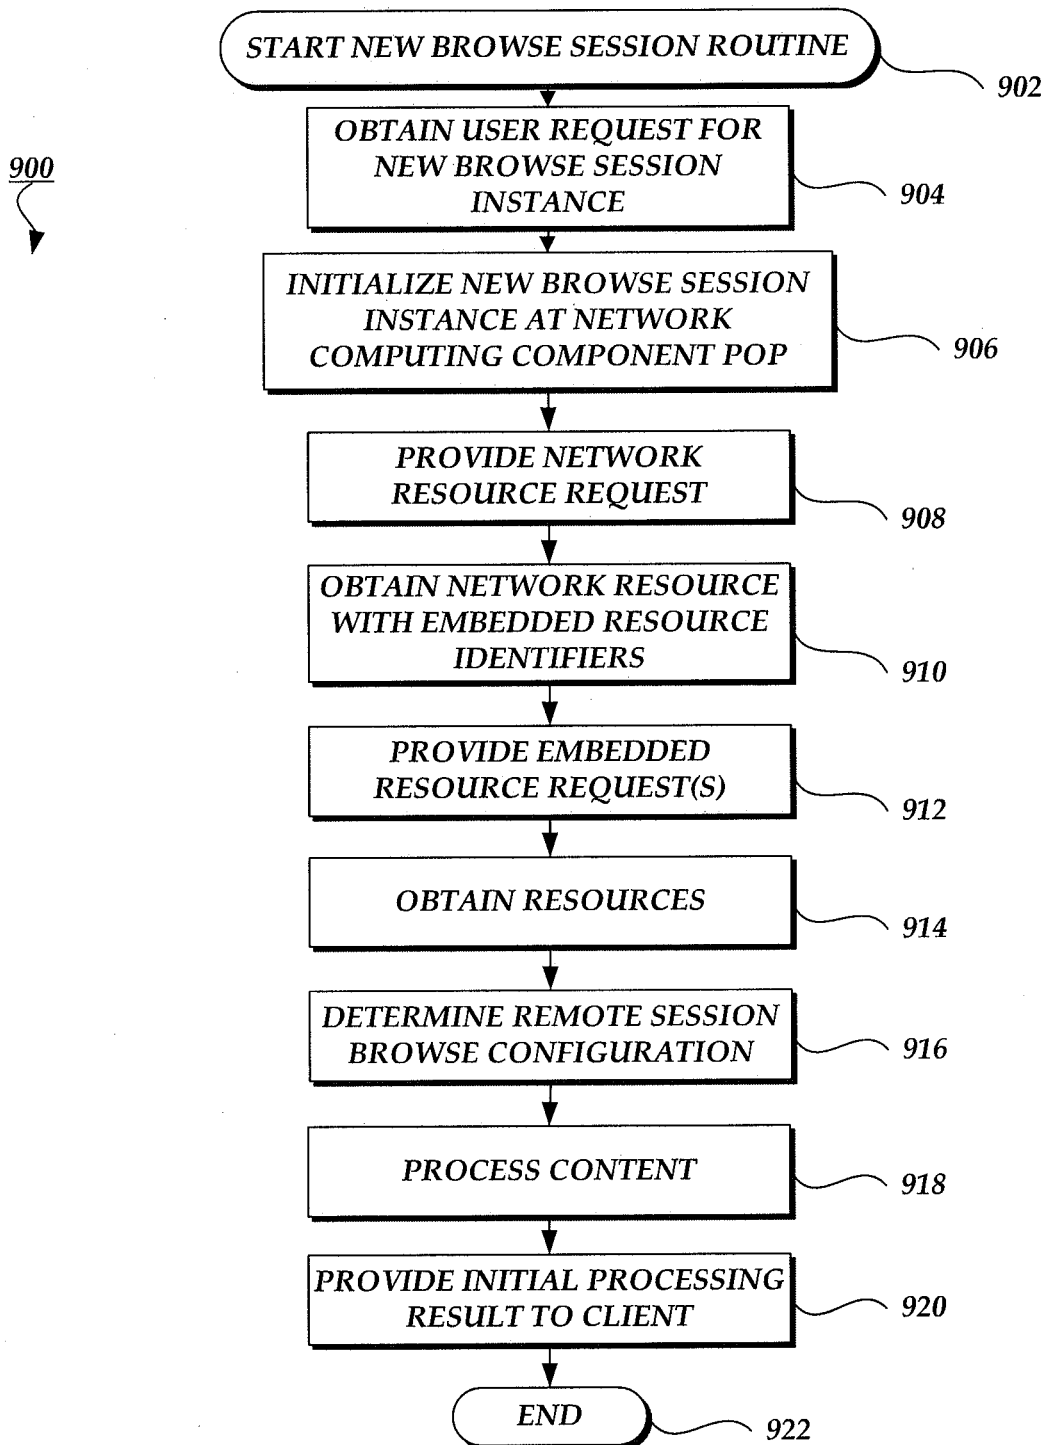


Fig. 7B.



**Fig. 8.**

**Fig. 9.**

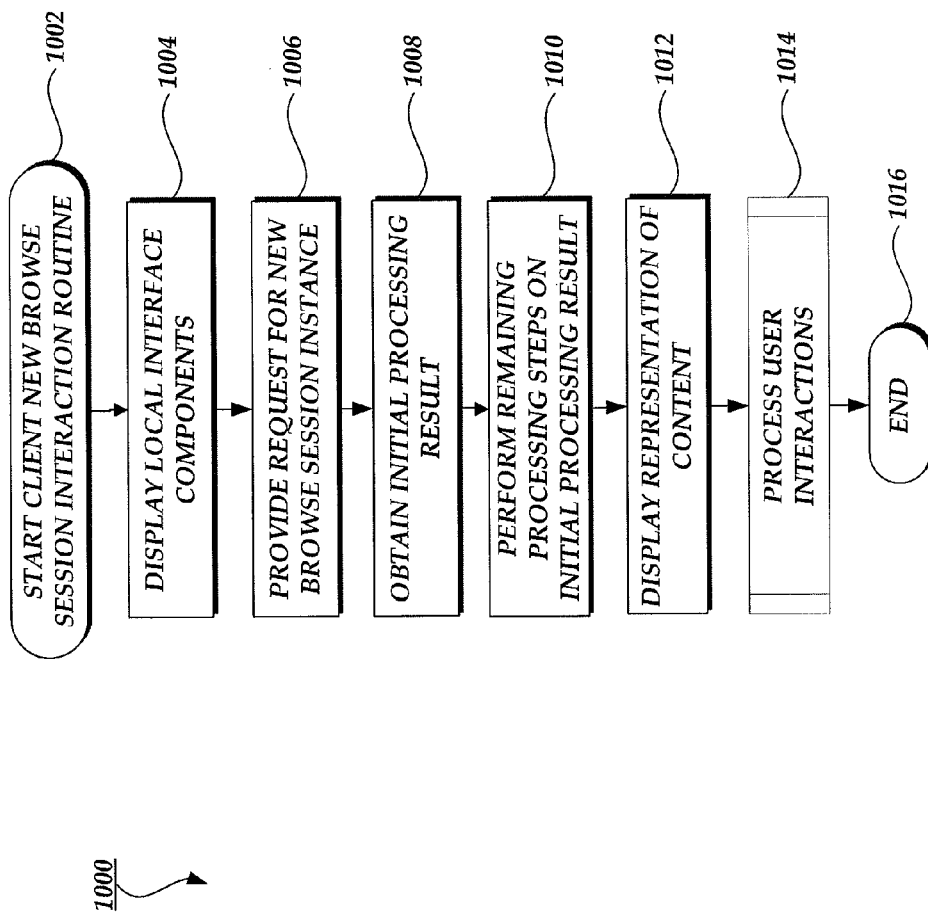
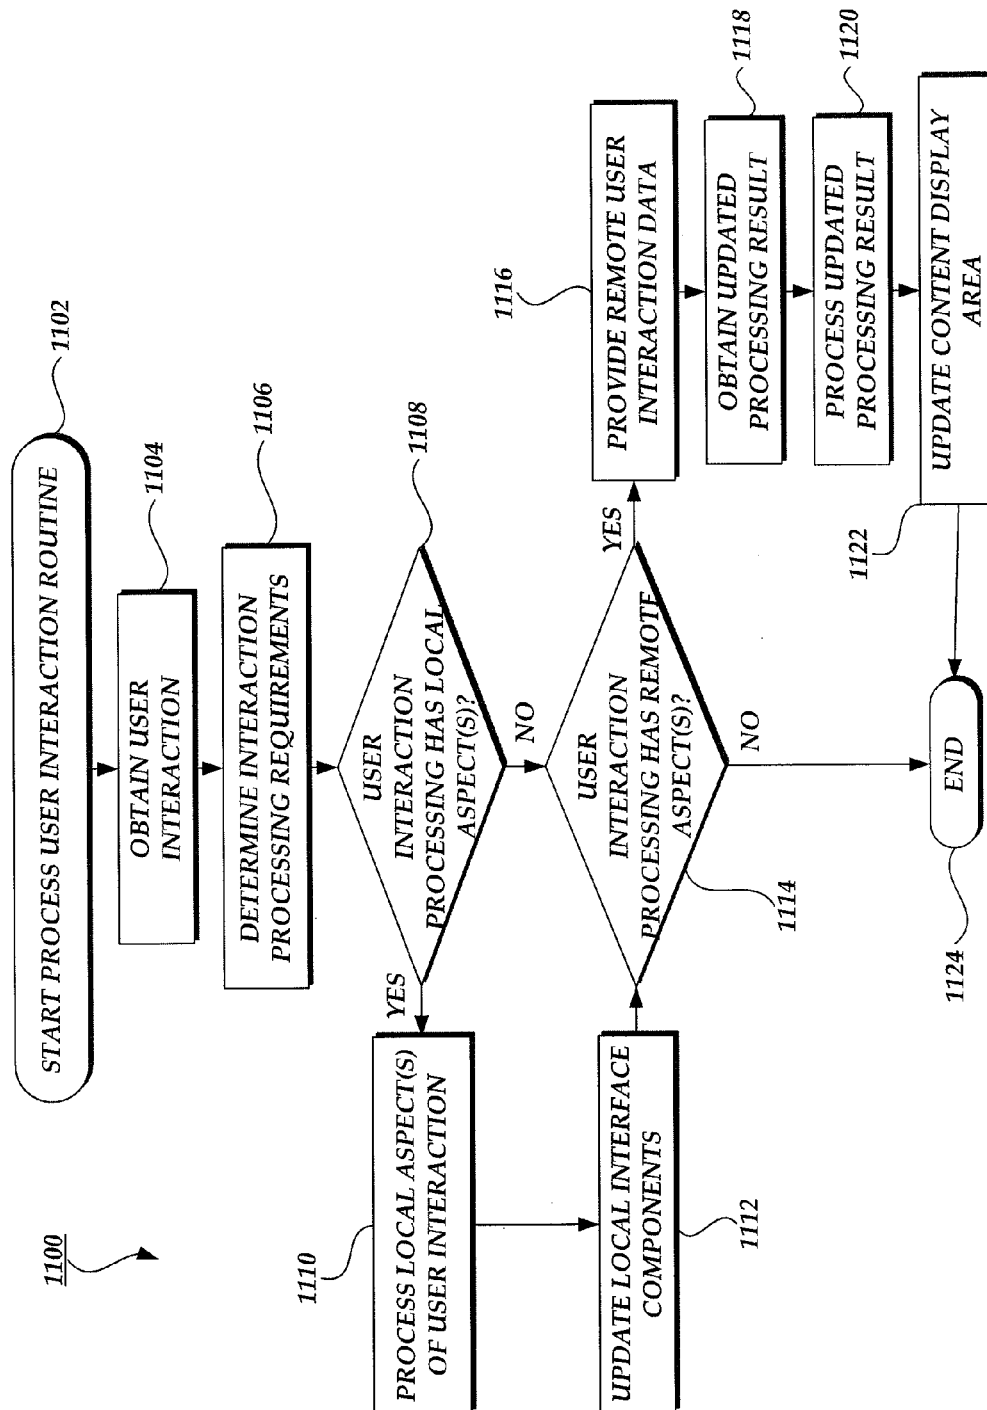
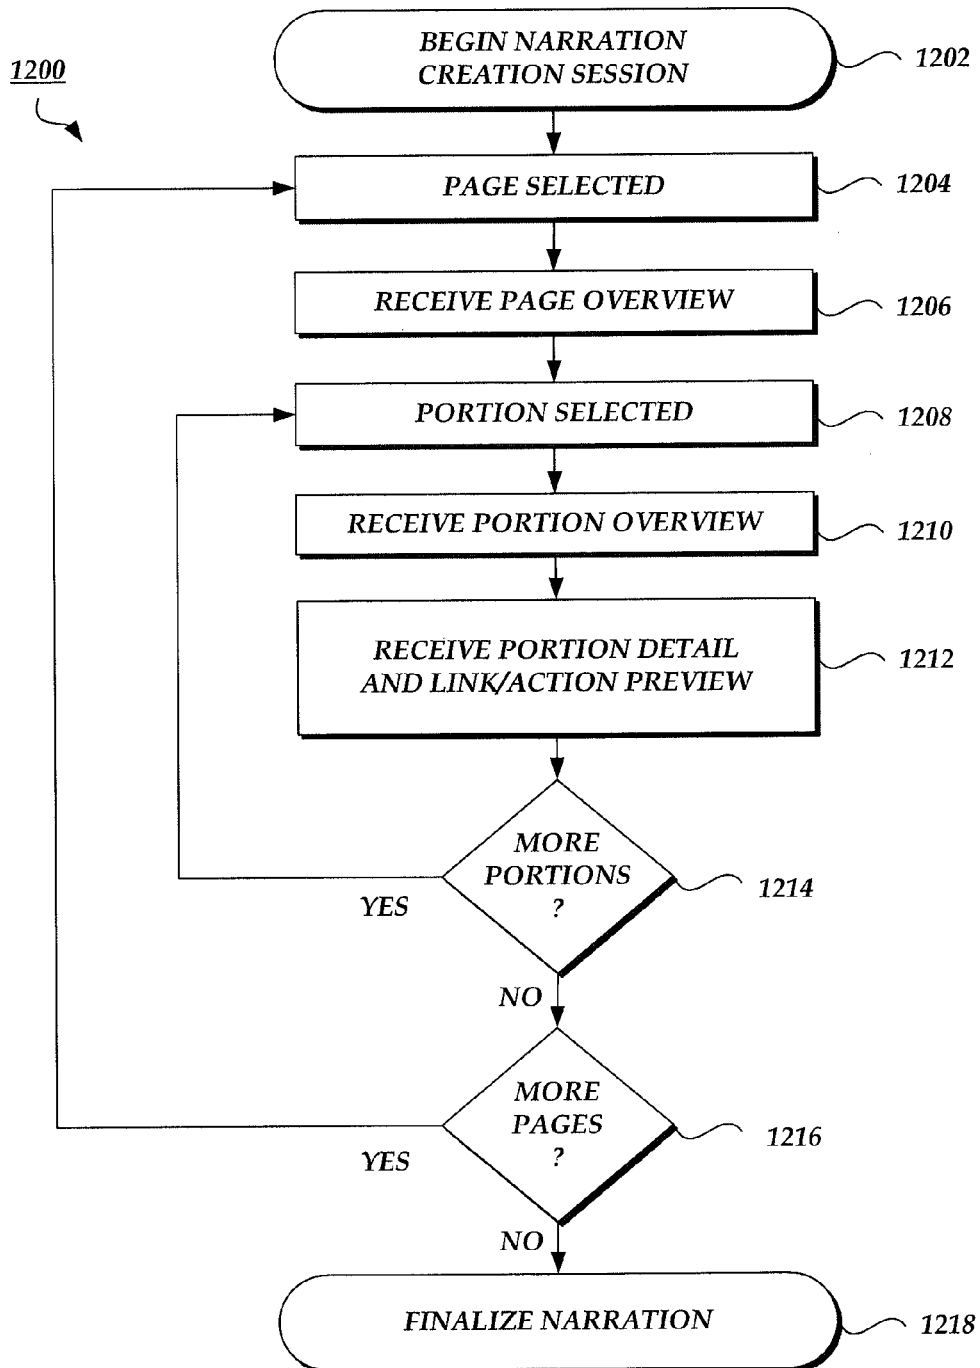
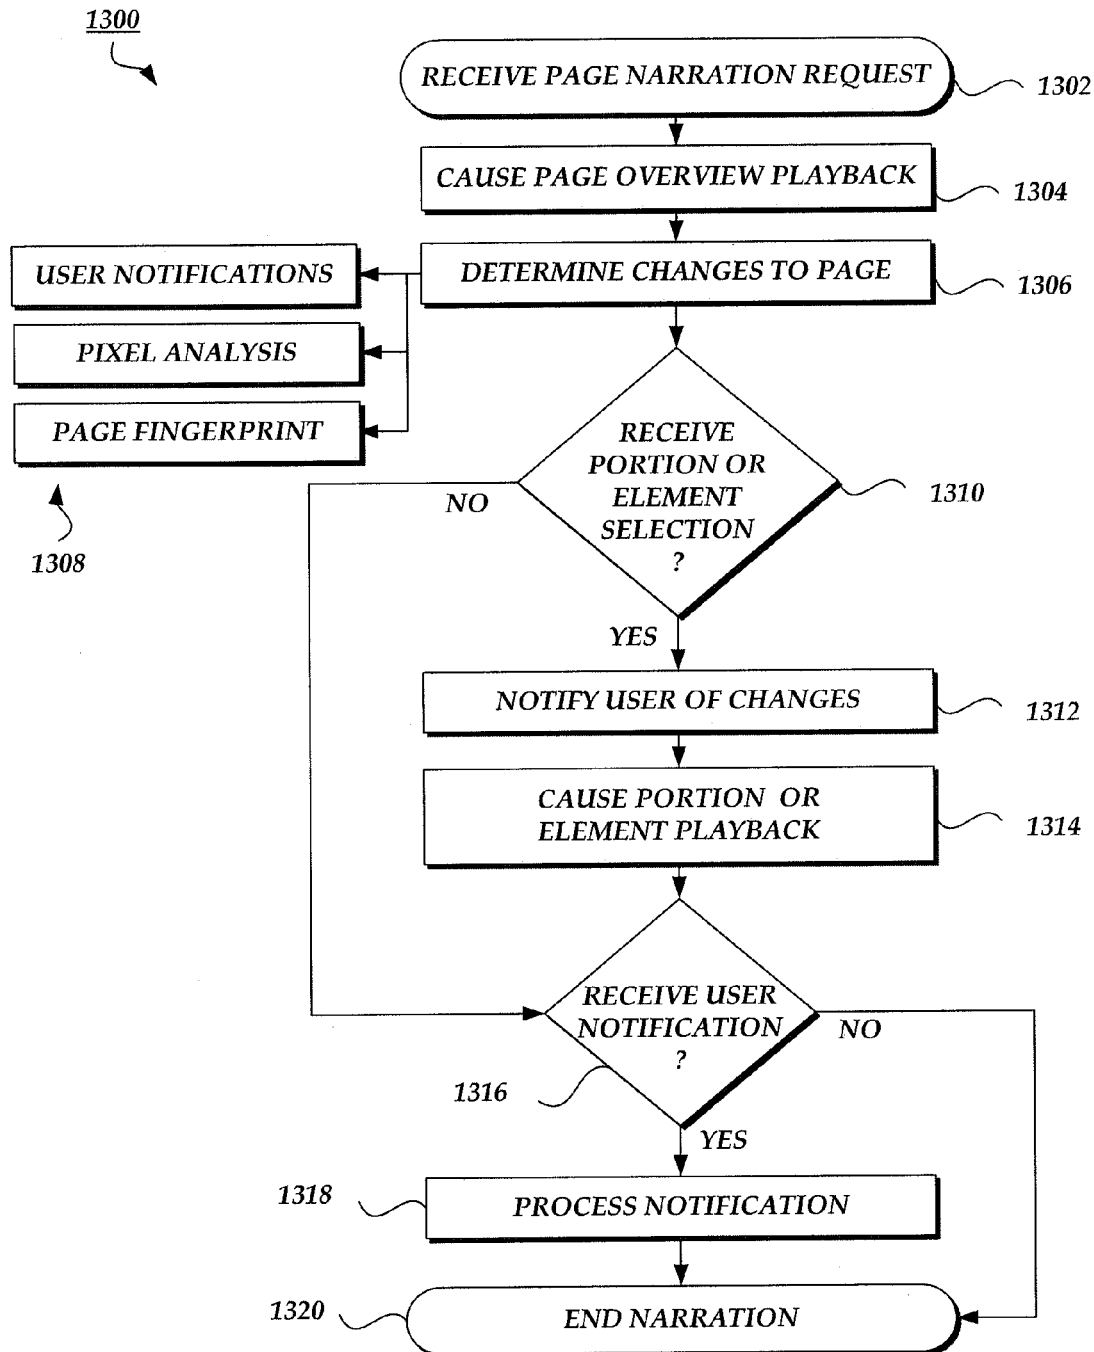
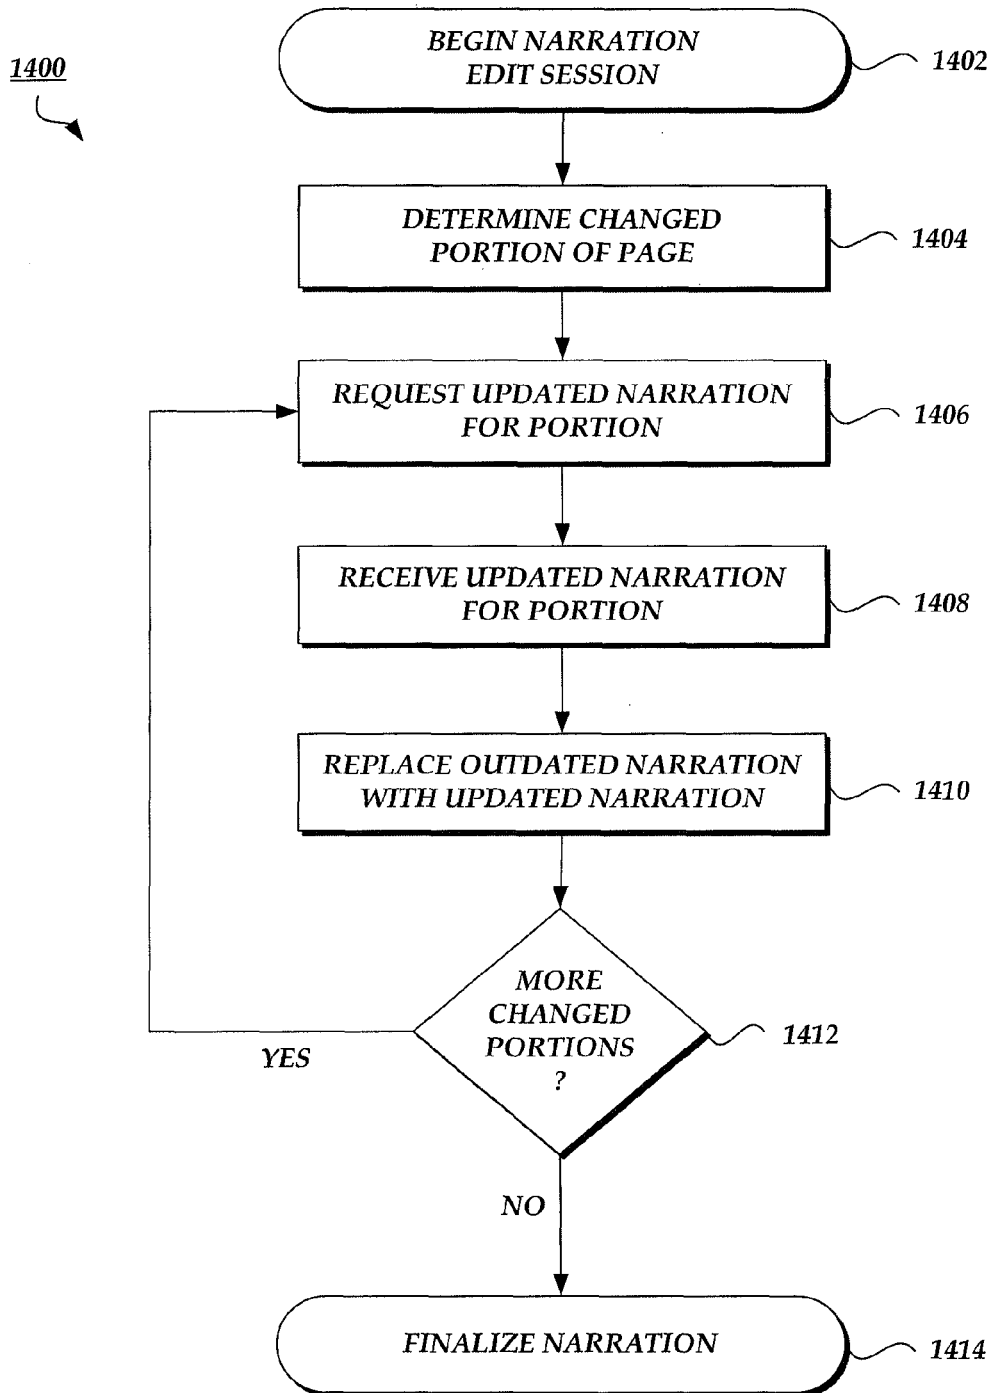


Fig. 10.

*Fig. 11.*

**Fig.12.**

*Fig.13.*

*Fig.14.*

1

NARRATION OF NETWORK CONTENT

BACKGROUND

Generally described, computing devices and communication networks can be utilized to exchange information. In a common application, a computing device can request content from another computing device via the communication network. For example, a user at a personal computing device can utilize a software browser application, typically referred to as a browser software application, to request a Web page from a server computing device via the Internet. In such embodiments, the user computing device can be referred to as a client computing device and the server computing device can be referred to as a content provider.

With reference to an illustrative example, a requested Web page, or original content, may be associated with a number of additional resources, such as images or videos, that are to be displayed with the Web page. In one specific embodiment, the additional resources of the Web page are identified by a number of embedded resource identifiers, such as uniform resource locators ("URLs"). In turn, software on the client computing devices, such as a browser software application, typically processes embedded resource identifiers to generate requests for the content. Accordingly, in order to satisfy a content request, one or more content providers will generally provide client computing devices data associated with the Web page as well as the data associated with the embedded resources.

Once the client computing device obtains the Web page and associated additional resources, the content may be processed in a number of stages by the software browser application or other client computing device interface. For example, and with reference to the above illustration, the software browser application may parse the Web page to process various HTML layout information and references to associated resources, may identify and process Cascading Style Sheets ("CSS") information, may process and instantiate various Javascript code associated with the Web page, may construct a native object model to represent one or more components of the Web page, and may calculate various layout and display properties of the processed content for presentation to a user.

Generally described, Web pages are primarily visual, consisting of textual elements, images, videos, animation, etc. Accordingly, users with visual impairments, an inability to read, or limited visual displays may not be able to consume and interact with the visual content on the page. In similar manner, Web pages contain aural elements, such as music, sound effects, etc. Such aural elements may be difficult to consume and interact with by users with hearing impairments or limited sound resources or restrictions on the use of sound resources.

BRIEF DESCRIPTION OF THE DRAWINGS

The foregoing aspects and many of the attendant advantages of this invention will become more readily appreciated as the same become better understood by reference to the following detailed description, when taken in conjunction with the accompanying drawings, wherein:

FIG. 1 is a block diagram illustrative of a content delivery environment including a number of client computing devices, content provider, a content delivery network service provider, and a network computing provider;

FIG. 2 is a block diagram of the content delivery environment of FIG. 1 illustrating the generation and processing of a

2

new browse session request from a client computing device to a network computing provider;

FIG. 3 is a block diagram of the content delivery environment of FIG. 1 illustrating the generation and processing of a request for a network resource from a network computing provider to a content provider;

FIG. 4 is a block diagram of the content delivery environment of FIG. 1 illustrating the generation and processing of one or more requests corresponding to one or more embedded resources from a network computing provider to a content provider and content delivery network;

FIG. 5 is a block diagram of the content delivery environment of FIG. 1 illustrating the generation and processing of browse session data and user interaction data between a network computing provider and client computing device;

FIG. 6 is a block diagram of the content delivery environment of FIG. 1 illustrating the generation and processing of an additional new browse session request from a client computing device to a network computing provider;

FIG. 7A is a user interface diagram depicting an illustrative browser interface and display of browse session content and narration features;

FIG. 7B is a user interface diagram depicting an illustrative browser interface and display of browse session content and playback features;

FIG. 7C is a user interface diagram depicting an illustrative browser interface and display of browse session content and alternative playback features;

FIG. 8 is a diagram depicting illustrative browser content processing actions as a series of processing subsystems;

FIG. 9 is a flow diagram illustrative of a new browse session routine implemented by network computing provider;

FIG. 10 is a flow diagram illustrative of a client new browse session interaction routine implemented by a client computing device;

FIG. 11 is a flow diagram illustrative of a process user interaction routine implemented by a client computing device;

FIG. 12 is a flow diagram illustrative of a routine implemented by a network computing provider for receiving and processing narration of a network resource;

FIG. 13 is a flow diagram illustrative of a routine implemented by a network computing provider for processing and responding to a request for playback of a narration; and

FIG. 14 is a flow diagram illustrative of a routine implemented by a network computing provider for processing changes to a narration of a network resource.

DETAILED DESCRIPTION

Generally described, the present disclosure is directed to the rendering of network resources as requested by a client computing device. Specifically, aspects of the disclosure will be described with regard to the creation, modification, and playback of aural or visual narrations of content associated with network resources, such as web pages. For example, a narration may be a spoken or computer-generated aural recitation of the textual elements of the content, an aural description of visual elements of the content, a textual transcription of any aural elements in the content, visual effects highlighting various portions of the content, commentary about the content or related topics, etc. Additionally, other aspects of the present disclosure relate to the storage and maintenance of narration information. Still further, aspects of the present disclosure relate to modification of the narrations in response to changes in the corresponding network resources. In still

another aspect of the present disclosure, client computing devices can be provisioned with narration information. Although aspects of the embodiments described in the disclosure will focus, for the purpose of illustration, on the processing by a network computing provider, one skilled in the art will appreciate that the techniques disclosed herein may be applied to any number of hardware or software processes or applications. Further, although various aspects of the disclosure will be described with regard to illustrative examples and embodiments, one skilled in the art will appreciate that the disclosed embodiments and examples should not be construed as limiting.

With reference to an illustrative example, a user may cause a client computing device to load a software browser application (henceforth referred to as a “browser”) for accessing content provided by one or more content providers. Illustratively, the accessed content may include a collection of one or more network resources (e.g., a Web page) and embedded resources, such as images, video, audio, text, executable code, and other resources. In one embodiment, the browser may have a content display area or pane for displaying the accessed network content in addition to one or more local interface components, such as toolbars, menus, buttons, or other user interface controls. Local interface components may be managed and controlled by the software browser application or any other process executing or implemented locally at the client computing device. Illustratively, managing user interface controls locally may allow for a responsive user interface, as interactions by the user are processed locally on the client computing device.

Subsequent to the browser being loaded, a user or automated browser process may cause the client computing device to transmit a request to access content from a content provider. For example, the user may select or enter a URL, (e.g., <http://www.xyzwebsite.com>) into a browser window, causing the client computing device to transmit a request for content corresponding to the URL. In an illustrative embodiment, the request can be transmitted to a network computing provider. Responsive to the request received from the client computing device, the network computing provider may instantiate or cause to have instantiated one or more computing components associated with the network computing provider that will host a browser software application. For example, the network computing provider can instantiate, or cause to have instantiated, an instance of a virtual machine that includes a software browser application capable of requesting resources from a communication network. Illustratively, in some situations, one or more devices associated with the network computing provider may be located in a data center or other robustly networked computing environment, and, as compared to the client computing device, may experience relatively little latency or delay when obtaining network resources.

Using the instantiated network computing components, the network computing provider may request the identified network resource(s) from one or more content providers, a content delivery network, or a local or associated cache component. For example, the browser software application on the instantiated network computing component can process a primary network resource and then generate additional content requests for content identified in one or more embedded resource identifiers (e.g. pictures, video files, etc.). Illustratively, in the case of other, non-browser, applications, network resources, or content may include any file type or format known in the art and supported by the specific software application.

Having obtained the requested content (e.g., the requested network resource and embedded resources), the network computing provider may identify one or more browsing configurations for the requested content, based on the processing requirements for one or more portions of the requested content. In addition, the network computing provider may determine that a narration of the requested content is available. The narration may have been received by the network computing provider at a previous time from a separate user of a client computing device. The narration may be a spoken recitation of the textual elements of the content, a spoken description of visual elements of the content, a textual transcription of any sound in the content, visual effects highlighting various portions of the content, etc. The network computing provider may provide access to the narration in addition to performing other processing of the network resource according to the browsing configuration.

In an illustrative embodiment, the network computing provider can analyze the requested content and compare the results to the requested network resource as it existed when the narration was created. The network computing provider can modify the narration if it determines that the requested content has changed since the narration was created. For example, the network computing provider can remove narration of content that no longer exists as part of the network resource, insert a warning to the user that the narration may not be accurate, etc. The network computing provider can provide the narration to the user of the client computing device all at once, as in streaming content with audio, in response to the selection of certain portions of the requested content at the client computing device, etc. The network computing provider can also cause elements to be displayed with the requested content to indicate to the user that a narration is available. The user of the client computing device can continue to interact with the network resource in a manner similar or identical to the manner of interaction if there were no narration.

FIG. 1 is a block diagram illustrative of a networked computing environment 100 for the management and processing of content requests. As illustrated in FIG. 1, the networked computing environment 100 includes a number of client computing devices 102 (generally referred to as clients) for requesting content and content processing from a content provider 104, CDN service provider 106, or network computing provider 107. In an illustrative embodiment, the client computing devices 102 can correspond to a wide variety of computing devices including personal computing devices, laptop computing devices, hand-held computing devices, terminal computing devices, mobile devices (e.g., mobile phones, tablet computing devices, etc.), wireless devices, various electronic devices and appliances and the like. In an illustrative embodiment, the client computing devices 102 include necessary hardware and software components for establishing communications over a communication network 108, such as a wide area network or local area network. For example, the client computing devices 102 may be equipped with networking equipment and browser software applications that facilitate communications via the Internet or an intranet. The client computing devices 102 may have varied local computing resources, such as central processing units and architectures, memory, mass storage, graphics processing units, communication network availability and bandwidth, etc.

The networked computing environment 100 can also include a content provider 104 in communication with the one or more client computing devices 102 or other service providers (e.g., CDN service provider 106, network computing

5

provider 107, etc.) via the communication network 108. The content provider 104 illustrated in FIG. 1 corresponds to a logical association of one or more computing devices associated with a content provider. Specifically, the content provider 104 can include a web server component 110 corresponding to one or more server computing devices for obtaining and processing requests for content (such as Web pages) from the client computing devices 102 or other service providers. The content provider 104 can further include an origin server component 112 and associated storage component 114 corresponding to one or more computing devices for obtaining and processing requests for network resources from the CDN service provider. The content provider 104 can still further include an application server computing device 111, such as a data streaming server, for processing streaming content requests. One skilled in the relevant art will appreciate that the content provider 104 can be associated with various additional computing resources, such as additional computing devices for administration of content and resources, DNS name servers, and the like. For example, although not illustrated in FIG. 1, the content provider 104 can be associated with one or more DNS name server components that would be authoritative to resolve client computing device DNS queries corresponding to a domain of the content provider.

With continued reference to FIG. 1, the networked computing environment 100 can further include a CDN service provider 106 in communication with the one or more client computing devices 102 and other service providers via the communication network 108. The CDN service provider 106 illustrated in FIG. 1 corresponds to a logical association of one or more computing devices associated with a CDN service provider. Specifically, the CDN service provider 106 can include a number of Point of Presence ("POP") locations 116, 122, 128 that correspond to nodes on the communication network 108. Each CDN POP 116, 122, 128 includes a DNS component 118, 124, 130 made up of a number of DNS server computing devices for resolving DNS queries from the client computers 102. Each CDN POP 116, 122, 128 also includes a resource cache component 120, 126, 132 made up of a number of cache server computing devices for storing resources from content providers and transmitting various requested resources to various client computers. The DNS components 118, 124, and 130 and the resource cache components 120, 126, 132 may further include additional software and/or hardware components that facilitate communications including, but not limited to, load balancing or load sharing software/hardware components.

In an illustrative embodiment, the DNS component 118, 124, 130 and resource cache component 120, 126, 132 are considered to be logically grouped, regardless of whether the components, or portions of the components, are physically separate. Additionally, although the CDN POPs 116, 122, 128 are illustrated in FIG. 1 as logically associated with the CDN service provider 106, the CDN POPs will be geographically distributed throughout the communication network 108 in a manner to best serve various demographics of client computing devices 102. Additionally, one skilled in the relevant art will appreciate that the CDN service provider 106 can be associated with various additional computing resources, such as additional computing devices for administration of content and resources, and the like.

With further continued reference to FIG. 1, the networked computing environment 100 can also include a network computing provider 107 in communication with the one or more client computing devices 102, the CDN service provider 106, and the content provider 104 via the communication network

6

108. The network computing provider 107 illustrated in FIG. 1 also corresponds to a logical association of one or more computing devices associated with a network computing provider. Specifically, the network computing provider 107 can include a number of Point of Presence ("POP") locations 134, 142, 148 that correspond to nodes on the communication network 108. Each POP 134, 142, 148 includes a network computing component (NCC) 136, 144, 150 for hosting applications, such as data streaming applications, via a number of instances of a virtual machine, generally referred to as an instance of an NCC. One skilled in the relevant art will appreciate that NCC 136, 144, 150 would include physical computing device resources and software to provide the multiple instances of a virtual machine or to dynamically cause the creation of instances of a virtual machine. Such creation can be based on a specific request, such as from a client computing device, or the NCC can initiate dynamic creation of an instance of a virtual machine on its own. Each NCC POP 134, 142, 148 also includes a storage component 140, 146, 152 made up of a number of storage devices for storing any type of data used in the delivery and processing of network or computing resources, including but not limited to user data, state information, processing requirements, historical usage data, and resources from content providers that will be processed by an instance of an NCC 136, 144, 150 and transmitted to various client computers, etc. The NCCs 136, 144, 150 and the storage components 140, 146, 152 may further include additional software and/or hardware components that facilitate communications including, but not limited to, load balancing or load sharing software/hardware components for selecting instances of a virtual machine supporting a requested application and/or providing information to a DNS nameserver to facilitate request routing.

In an illustrative embodiment, NCCs 136, 144, 150 and the storage components 140, 146, 152 are considered to be logically grouped, regardless of whether the components, or portions of the components, are physically separate. For example, a network computing provider 107 may maintain separate POPs for providing the NCC and the storage components. Additionally, although the NCC POPs 134, 142, 148 are illustrated in FIG. 1 as logically associated with a network computing provider 107, the NCC POPs will be geographically distributed throughout the communication network 108 in a manner to best serve various demographics of client computing devices 102. Additionally, one skilled in the relevant art will appreciate that the network computing provider 107 can be associated with various additional computing resources, such as additional computing devices for administration of content and resources, and the like. Even further, one skilled in the relevant art will appreciate that the components of the network computing provider 107 and components of the CDN service provider 106 can be managed by the same or different entities.

With reference now to FIGS. 2-6, the interaction between various components of the networked computing environment 100 of FIG. 1 will be illustrated. Specifically, FIGS. 2-6 illustrate the interaction between various components of the networked computing environment 100 for the exchange of content between a client computing device 102 and a content provider 104 via the network computing provider 107. For purposes of the example, the illustrations have been simplified such that many of the components utilized to facilitate communications are not shown. One skilled in the relevant art will appreciate that such components can be utilized and that additional interactions would accordingly occur without departing from the spirit and scope of the present disclosure.

With reference to FIG. 2, the process can begin with the generation and processing of a browse session request from a client computing device 102 to a network computing provider 107 will be described. Illustratively, the client computing device 102 may load a browser for viewing network content in response to an event or user request. Subsequent to the browser being loaded, the browser may be implemented to request a new browse session. From the perspective of the user of the client computing device, the request for the new browse session corresponds to the intended request to transmit the request to one or more corresponding content providers 104. Illustratively, this request may be generated automatically as a result of the browser loading (e.g., a request for a default or “home” page), or may be generated as a result of a user following a link or entering a network address into an address bar. As illustrated in FIG. 2, the browse session request is transmitted first to a network computing provider 107. In an illustrative embodiment, the network computing provider 107 utilizes a registration application program interface (“API”) to accept browse session requests from the client computing device 102. The browse session request can include network address information corresponding to a requested network resource, which may be in any form, including, but not limited to, an Internet Protocol (“IP”) address, a URL, a Media Access Control (“MAC”) address, etc.

Subsequent to the receipt of the browse session request, the network computing provider 107 may select an associated network computing component (hereinafter “NCC”) point of presence (hereinafter “POP”) such as NCC POP 142 to service the browse session request. The selection of the NCC POP may determine the processing and network resources available to the instantiated virtual machine. The selection of processing and network resources and the provisioning of software at the NCC POP instance may be done, at least in part, in order to optimize communication with content providers 104 and client computing devices 102. As illustrated in FIG. 2, the identification data received with the browse session request can be passed to the NCC POP instance servicing the browse session request.

With reference to FIG. 3, an illustrative interaction for generation and processing of a request for a network resource from a network computing provider 107 to a content provider 104 will be described. As illustrated in FIG. 3, the selected NCC POP 142 may generate a browse session corresponding to one or more content providers based on a browse session request, such as the illustrative browse session request depicted in FIG. 2 above. Illustratively, instantiating a new browse session instance may include loading a new virtual machine instance and/or browser instance at the NCC POP 142, reserving or allocating device memory, storage or cache space, processor time, network bandwidth, or other computational or network resources for the new browse session.

Subsequent to initializing a new browse session instance, NCC POP 142 may provide a request for a network resource to a content provider 104 based on a network address included in the browse session request. For example, a browse session request may include a URL for a Web page, such as “http://www.xyzsite.com/default.htm.” NCC POP 142 may resolve the URL to an IP address through a DNS resolver associated with the network computing provider (not shown), and may request the Web page from the content provider 104 at the resolved IP address. In various embodiments, a network resource may be retrieved from any combination of content providers, content delivery network (hereinafter “CDN”) servers, or caches associated with the network computing provider 107. For example, the network computing provider

may check if a resource is stored in a local cache or in another server or service provider associated with the network computing provider 107. If a network resource is stored in a local or associated location, the NCC POP 142 may retrieve the network resource from the local or associated location rather than from the third party content provider 104 or CDN service provider 106. Illustratively, the NCC POP 142 may provide requests for any number of network resources as included in the browse session request, and may obtain these network resources from any number of different sources, sequentially or in parallel.

As illustrated in FIG. 3, the content provider 104 receives the resource request from the NCC POP 142 and processes the request accordingly. In one embodiment, the content provider 104 processes the resource request as if it were originally provided by the client computing device 102. For example, the content provider 104 may select the type of content, ordering of content, or version of content according to the requirements of the requesting client computing device 102. In another embodiment, the content provider 104 may be provided with information that provides information associated with the NCC POP 142 for utilization in providing the requested content (e.g., an available amount of processing resources or network bandwidth).

Subsequent to obtaining the requested network resource from the content provider 104 (or other source designated by the content provider), the NCC POP 142 may process the network resource to extract embedded resource identifiers and gather information for determination of a remote session browsing configuration. For example, a network resource such as a Web page may include embedded CSS style information and Javascript as well as embedded resource identifiers to additional resources, such as text, images, video, audio, animation, executable code, and other HTML, CSS, and Javascript files. In the process of extracting the embedded resource identifiers, the NCC POP 142 may gather information about the processed network resources for later use in the determination of a remote session browsing configuration as discussed below with reference to FIG. 4.

With reference to FIG. 4, an illustrative interaction for generation and processing of one or more requests corresponding to one or more embedded resources from a network computing provider to a content provider and content delivery network is disclosed. As illustrated in FIG. 4, the selected NCC POP 142 may provide resource requests to one or more sources of content such as content provider 104 and CDN POP 116. The resource requests may correspond to embedded resources based on one or more embedded resource identifiers extracted from a requested network resource (e.g., a Web page) as described in FIG. 3 above. In various embodiments, embedded resources may be retrieved from any combination of content providers, CDN servers, or caches associated with the network computing provider 107. For example, the network computing provider may check if an embedded resource is stored in a local cache or in another server or service provider associated with the network computing provider 107. If an embedded resource is stored in a local or associated location, the NCC POP 142 may retrieve the embedded resource from the local or associated location rather than the third party content provider or CDN. Illustratively, the NCC POP 142 may provide requests for any number of embedded resources referenced by a network resource, and may obtain these embedded resources from any number of different sources, sequentially or in parallel. Subsequent to obtaining the requested resources, the NCC POP 142 may process the resources and requested content to determine a

remote session browsing configuration for the processing and communication of content to the client computing device **102**.

With reference to FIG. 5, an illustrative interaction for generation and processing of processing results and user interaction data between a network computing provider and client computing device is disclosed. As previously described, in one embodiment, the respective browsers on the instantiated network computing component and the client computing device **102** can exchange browsers' session information related to the allocation and processing of the requested resources at the instantiated network computing component and client computing device. As illustrated in FIG. 5, the selected NCC POP **142** may provide an initial processing result to the client computing device **102** over the network **108**. The initial processing result may correspond to requested network content, such as a Web page, along with associated embedded resources processed by the NCC POP **142** in accordance with a selected remote session browsing configuration as described in FIG. 4 above. The NCC POP **142** also makes a determination of which additional processes will be conducted at the NCC POP **142**, at the client computing device **102**, or both, as illustrated in FIG. 5. Subsequent to receiving an initial processing result and the allocation of processes, the client computing device **102** may perform any remaining processing actions on the initial processing result as required by the selected remote session browsing configuration, and may display the fully processed content in a content display area of a browser. The client computing device **102** may process any local user interactions with local interface components or content elements locally, and may provide user interactions requiring remote processing to the network computing provider **107**. The network computing provider **107** may provide updated processing results to the client computing device **102** in response to changes to the content or remote user interaction data from the client computing device **102**.

With reference to FIG. 6, a block diagram of the content delivery environment of FIG. 1 illustrating the generation and processing of an additional new browse session request from a client computing device to a network computing provider is disclosed. As illustrated in FIG. 6, a second new browse session request may be sent to network computing provider **107** from client computing device **102** across network **108**. In an illustrative embodiment, the network computing provider **107** utilizes a registration API to accept browse session requests from the client computing device **102**.

The additional browse session request may be generated by a client computing device **102** in response to a user opening up a new browser window with a new content display area, opening a new content display area in an existing browser window (e.g., opening a new tab in a browser), requesting new network content in an existing content display area (e.g., following a link to a new network resource, or entering a new network address into the browser), or any other user interaction. For example, a user browsing a first Web page corresponding to a first browse session instance may follow a link that opens a new tab or browser window to view a second Web page. In one embodiment, any required steps of obtaining and processing content associated with the second Web page may be performed by the currently instantiated network computing component in which the browser can handle the processing of both resource requests. In another embodiment, the client computing device **102** request may be processed as a new browse session request to the network computing provider **107**, including the network address of the second Web page. In this embodiment, the browser on the client comput-

ing device may not specifically request a separate browse session, and a user's interaction with the browser on the client computing device **102** may appear to be part of a same browsing session. As described above with regard to FIGS. 2 and 3, the network computing provider **107** may cause an instantiation of a network computing component for obtaining and processing content associated with the second web page. In other embodiments, a new browse session request may be generated by the client computing device **102** corresponding to sections of a network resource (e.g., frames of a Web page), individual network resources, or embedded resources themselves, data objects included in a set of content, or individual network resources.

Illustratively, the additional browse session request may include any number of pieces of data or information including, but not limited to, information associated with a user, information associated with the client computing device **102** (e.g., hardware or software information, a device physical or logical location, etc.), information associated with the network **108**, user or browser preferences (e.g., a requested remote session browse protocol, a preference list, a decision tree, or other information), information associated with the network computing provider **107**, information associated with one or more pieces of requested network content (e.g., the network address of a network resource), etc. Requested content may include any manner of digital content, including Web pages or other documents, text, images, video, audio, executable scripts or code, or any other type of digital resource.

As illustrated in FIG. 6, subsequent to the receipt of the browse session request the network computing provider **107** may select an associated network computing component such as NCC POP **142** to service the browse session request. As discussed above with reference to FIG. 2, a network computing provider **107** may select an NCC POP to service a browse session request based on any number of factors, including, but not limited to available NCC POP resources (e.g., available memory, processor load, network load, etc), a financial cost of servicing the browse session request at the NCC POP, the NCC POP location respective to a client computing device **102**, content provider **112**, or CDN POP **116**, a NCC POP cache status (e.g., whether a requested resource is already stored in an NCC POP cache), etc. In one embodiment, the network computing provider **107** may select a number of NCC POPs to service a browse session request. Illustratively, although the network computing provider **107** is depicted here for purposes of illustration as selecting NCC POP **142**, the network computing provider **107** may select any extant NCC POP to service the browse session request. For example, a single client computing device **102** may simultaneously or sequentially provide three different browse session requests to the network computing provider **107** corresponding to different network resources. The network computing provider **107** may select different NCC POPs for each browse session request, the same NCC POP for all three browse session requests, or any combination thereof. As discussed above, the decision whether to select a different NCC POP than was utilized for a previous or simultaneous browse session request may be made on the basis of available system resources, randomly, or according to any other factor as discussed above and with regards to FIG. 2.

FIG. 7A is a user interface diagram depicting an illustrative browser interface and display of browse session content. As described above with reference to FIG. 5, a browser **700** may have a content display area **702**, as well as one or more one or more local interface components. These local interface components may include toolbars, menus, buttons, address bars,

11

scroll bars, window resize controls, or any other user interface controls. For example, the local interface components can include a toolbar with a narrate button **704** to facilitate the creation of narrations for content displayed in the content display area **702**. Illustratively, local interface components may be displayed as separate from the content display area **702** or may be overlaid or embedded in the content display area **702**.

Interactions with local interface components may be treated as local user interactions or remote user interactions depending on the processing required by the interaction and the remote session browsing configuration. For example, the selection of a preferences option in a browser menu may be handled entirely as a local user interaction by a browser. The processing required to display the menu, provide visual feedback regarding the selection, display the preferences window, and process the changes made to the browser preferences may be performed locally. As discussed above, processing user interactions locally may provide greater responsiveness at the browser as opposed to sending user interaction data to the NCC POP **142** for processing. As another example, when using a remote session browsing configuration that specifies extensive processing on the NCC POP **142** (e.g., a remote session browsing configuration using a remote session communication protocol such as RDP), the selection of a content refresh button in a browser toolbar may be handled both as a local user interaction and a remote user interaction. The limited processing required to provide interface feedback corresponding to the button selection may be handled at the client computing device **102** in order to provide the appearance of interface responsiveness, while the refresh command, which may require processing of the network content displayed in the content display area of the browser, may be sent as user interaction data to the NCC POP **142** for processing. The NCC POP **142** may then transmit updated processing results corresponding to the refreshed network content back to the client computing device **102** for display.

In response to a user selecting the narration button **704**, the browser may enter a narration creation mode. During the narration creation mode, the user may verbally describe the content visible on the screen. Additionally, the user can select content that will be automatically narrated, such as by an automated text to speech software application. In some embodiments, the narration may include visual components, such as text and animation, in addition to or instead of a verbal description. While in narration creation mode, the normal cursor may be altered such that the user has visual feedback that the browser is in narration creation mode. For example, after selecting the narration button **704** to enter narration creation mode, the cursor may be changed to the cursor **706** illustrated in FIG. 7A.

The user may describe an image, such as image **710**. In one embodiment, the user can manipulate the cursor **706** to select a portion **708** of the page to describe. In FIG. 7A, the portion **708** contains the image **710** that the user will describe. The user can select other portions of the page and describe the visible contents of the other portions, in some cases reading the text on the screen, highlighting portions of text with the cursor, etc. In the case of a web page with hyperlinks and other interactive controls, the user can select a hyperlink and describe the destination referenced by the selected the hyperlink. In similar fashion, other interactive controls such as command buttons, radio buttons, and drop-down lists can be activated and described for the narration. The motion of the cursor on the screen can be recorded for playback, allowing the narrator to point to various elements when speaking about them. Various cursor, keyboard, menu, or spoken inputs can

12

also be recorded to display visual effects during playback, such as highlighting, etc. When the user has finished creating the narration, the user can again select the narration button **704** to stop narration creation and submit the narration to the network computing provider **107**. In some embodiments, the functions of the narration button **704** can be replaced by a menu option, keyboard combination, voice command, or any other input method. In similar fashion, the method of selecting portions of the content to narrate can include cursor actions, menu options, touch screen gestures, voice commands, electronic pen input, or any other input method.

FIG. 7B illustrates the browser interface of FIG. 7A in a narration playback mode. A user of a client computing device **102** can select the playback button **712** when a network resource, such as a web page, with an associated narration has been loaded. Because a consumer of a narration may be visually impaired, a method of receiving voice commands may be implemented so that the user can request web pages and narrations, activate hyperlinks, and perform other input verbally. The browser interface **700** may display visual indicators to the user when narration is available for a web page. For example, the title bar **714** may include an icon, the playback button **712** may normally be hidden and only made visible when a narration is available, etc.

During narration playback mode, some indication may be displayed on the browser interface **700** to inform the user which portion of the content is being described. For example, a highlight box **722** may appear around textual content being read to the user, or a link may be outlined **716** or otherwise highlighted to indicate to the user that a portion of the content is being generally described rather than read verbatim. In some embodiments, a user may be able to interact directly with the content as though there were no narration, even while narration playback is occurring. A user may position the mouse cursor **718** over a link and click the link, thereby loading a different page of content, even while narration playback is occurring with respect to the original content page. Some content may include dynamically changing portions, such as dynamic portion **720**. For example, a news web site may include a "top stories" section which updates on regular basis, sometimes in real time as events happen. The narration of such a portion may accordingly be out of date and inaccurate when played back at a later time. Therefore, in some embodiments a method may be included to determine whether a portion of the resource been altered since the narration was created.

FIG. 7C illustrates a browser interface **700** for interacting with narrated content. The interface can include a hovering activation link **724** for initiating playback of a narration, rather than the playback button **712** of FIG. 7B. Additionally, instead of the normal mouse cursor **718** of FIG. 7B, the interface of FIG. 7C can display a cursor **726** which indicates to the user that the browser is currently in narration playback mode. In some cases, the cursor **726** can be a visual component of the narration, and can indicate the position of the narrator's cursor on the screen during narration creation. FIG. 7C also illustrates an altered dynamic portion **720**, when compared to the browser of FIG. 7B. When the user selects the dynamic portion **720** for narration playback, whether by clicking, highlighting, etc., the browser **700** or network computing provider **107** can determine whether there has been a change in the content within the dynamic portion **720** since the narration was created, as described in detail below. In response to determining that there has been a change to the content within the selected dynamic portion **720**, the network computing provider **107** may provide an updated narration. For example, the network computing provider **107** may pro-

13

vide a more recently created narration for the dynamic portion 720, may provide a warning that the narration may be out of date, may remove the out of date portion from the narration, and the like.

FIG. 8 is a diagram depicting illustrative browser content processing actions as a series of processing subsystems 800. In many embodiments, a browser may process sets of content (e.g., network resources, such as web pages and associated embedded resources) in a series of processing actions. Illustratively, and as described above with reference to FIGS. 3-5, a remote session browsing configuration may specify a split between processing actions performed at a network computing provider (e.g., an NCC POP) and processing actions performed at a client computing device 102. This split may designate some processing actions to be performed by each of the NCC POP and client computing device 102, or may assign all processing actions to a single device or component. For example, an NCC POP may perform all of these various processing actions at the browse session instance, and send fully processed RDP processing results to the client computing device 102 for bitmap assembly and display. Any number of different remote session browsing configurations may be used by one or more browse sessions instances running at an NCC POP.

One of skill in the relevant art will appreciate that the subsystems shown here are depicted for the purpose of illustration, and are not intended to describe a necessary order or a definitive list of browser subsystems. Various browser software components may implement additional or fewer browser subsystems than are shown here, and may order the subsystems or corresponding processing actions in any number of different ways. Although the processing subsystems 800 depicted here for purposes of illustration are directed at the processing of Web pages or other Web content, one of skill in the relevant art will appreciate that the processing of other file types or network resources may be broken up in a similar manner. For example, one of skill in the relevant art will appreciate that a similar schema may be developed for the processing of images, video, audio, database information, 3 d design data, or any other file format or type of data known in the art. Similar schema may also be developed for any number of device operating system or software framework processing operations, such as scheduling, memory or file management, system resource management, process or service execution or management, etc. Further, although the HTML protocol and RDP remote session communication protocols are discussed herein for the purposes of example, one of skill in the relevant art will appreciate that a remote session browsing configuration may implement any number of remote communication protocols for any number of specified processing actions, and that a remote session browsing configuration may be formulated to perform any fraction or combination of the actions identified below at any combination of the client computing device 102 and network computing provider 107.

Illustratively, the first processing subsystem involved in the processing and display of network content is the networking subsystem 802. Illustratively, the networking subsystem 802 may be responsible for all communication between the browser and content provider, including local caching of Web content. The networking subsystem is generally limited by the performance of the user's network. A remote session browsing configuration that splits processing actions at the networking subsystem 802 might include a remote session browsing configuration utilizing an HTML remote session communication protocol, where one or more caching or

14

resource retrieval actions were performed at the NCC POP, but parsing and processing of the content was performed at the client computing device.

As network resources, such as HTML documents are downloaded from the server they may be passed to an HTML subsystem 804 which parses the document, initiates additional downloads in the networking subsystem, and creates a structural representation of the document. Modern browsers may also contain related subsystems which are used for XHTML, XML and SVG documents. A remote session browsing configuration that splits processing actions at the HTML subsystem 804 might include a remote session browsing configuration utilizing an HTML remote session communication protocol, where an initial HTML page is processed at the NCC POP in order to extract embedded resource identifiers, but additional parsing and processing of the content is performed at the client computing device. In another embodiment, a remote session browsing configuration that splits processing actions at the HTML subsystem 804 might perform initial processing to create the structural representation of the HTML document, and provides a processing result including the structural representation and associated embedded resources to the client computing device for processing.

When CSS is encountered, whether inside an HTML document or an embedded CSS document, it may be passed to a CSS subsystem 806 to parse the style information and create a structural representation that can be referenced later. Illustratively, a remote session browsing configuration that splits processing actions at a CSS subsystem 806 may construct a processing result including the CSS structural representation and HTML structural representation, and provide the processing result and associated embedded resources to the client computing device for processing.

HTML documents often contain metadata, for example the information described in a document header or the attributes applied to an element. The collections subsystem 808 may be responsible for storing and accessing this metadata. A remote session browsing configuration that splits processing actions at a collections subsystem 808 may construct a processing result including processed metadata along with any other structural representations discussed above, and provide the processing result and associated embedded resources to the client computing device for processing.

When Javascript is encountered, it may be passed directly to a JavaScript subsystem 810 responsible for executing the script. The Javascript subsystem 810 has been examined fully over the years, and may be one of the most well known browser subsystems in the art. A remote session browsing configuration that splits processing actions at a Javascript subsystem 810 may construct a processing result including an internal representation of one or more Javascript scripts, including, but not limited to state data or a representation of the script in a native or intermediate form, as well as any other processed structures or data discussed above, and provide the processing result and associated embedded resources to the client computing device for processing.

Because many JavaScript engines are not directly integrated into the browser, there may be a communication layer including the marshalling subsystem 812 between the browser and the script engine. Passing information through this communication layer may generally be referred to as marshaling. A remote session browsing configuration that splits processing actions at a marshalling subsystem 812 may construct a processing result including marshalling data as well as any other processed structures, scripts, or data dis-

15

cussed above, and provide the processing result and associated embedded resources to the client computing device for processing.

In some embodiments, JavaScript interacts with an underlying network resource such as a Web document through the Document Object Model APIs. These APIs may be provided through a native object model subsystem **814** that knows how to access and manipulate the document and is the primary interaction point between the script engine and the browser. Illustratively, a remote session browsing configuration that splits processing actions at a native object model subsystem **814** may construct a processing result including native object model state data or API calls as well as any other processed structures, scripts, or data discussed above, and provide the processing result and any other associated embedded resources to the client computing device for processing.

Once the document is constructed, the browser may need to apply style information before it can be displayed to the user. The formatting subsystem **816** takes the HTML document and applies styles. Illustratively, a remote session browsing configuration that splits processing actions at a formatting subsystem **816** may construct a processing result including an HTML representation with applied styles, as well as any other processed state data, API calls, structures, scripts, or data discussed above, and provide the processing result and any other associated embedded resources to the client computing device for processing.

In one embodiment, CSS is a block based layout system. After the document is styled, the next step, at a block building subsystem **818**, may be to construct rectangular blocks that will be displayed to the user. This process may determine things like the size of the blocks and may be tightly integrated with the next stage, layout. A remote session browsing configuration that splits processing actions at a block building subsystem **818** may construct a processing result including block information, as well as any other processed state data, API calls, structures, scripts, or data discussed above, and provide the processing result and any other associated embedded resources to the client computing device for processing.

Subsequent to the browser styling the content and constructing the blocks, it may go through the process of laying out the content. The layout subsystem **820** is responsible for this algorithmically complex process. Illustratively, a remote session browsing configuration that splits processing actions at a layout subsystem **820** may process the various state data, API calls, structures, scripts, or data discussed above to construct a processing result including layout information for the client computing device. Illustratively, an NCC POP may make use of various data or settings associated with the client computing device or browser (e.g., as provided in the initial browse session request) in order to generate a suitable layout for the client computing device. For example, a mobile device may provide a screen resolution and a display mode to the NCC POP. The NCC POP may base layout calculations on this screen resolution and display mode in order to generate a processing result corresponding to a content representation suitable for a browser running on the mobile device. Illustratively, in various embodiments, any other subsystem implemented by the NCC POP may make use of data associated with the client computing device or browser in generating a processing result for the client.

The final stage of the process may occur inside the display subsystem **822** where the final content is displayed to the user. This process is often referred to as drawing. A remote session browsing configuration that splits processing actions at the networking subsystem **802** might include a remote session browsing configuration utilizing an RDP remote session

16

communication protocol, where nearly all processing is performed at the NCC POP, and a processing result including bitmap data and low level interface data are passed to the client computing device for display.

FIG. **9** is a flow diagram illustrative of a new browse session routine **900** implemented by network computing provider **107** of FIG. **1**. New browse session routine **900** begins at block **902**. At block **904**, the network computing provider **107** receives a new browse session request from client computing device **102**. As previously described, the client computing device **102** may load a browser for viewing network content in response to an event or user request. Subsequent to the browser being loaded, the browser may be implemented request a new browse session. From the perspective of the user of the client computing device, the request for the new browse session corresponds to the intended request to transmit the request to one or more corresponding content providers **104**. Illustratively, this request may be generated automatically as a result of the browser loading (e.g., a request for a default or "home" page), or may be generated as a result of a user following a link or entering a network address into an address bar. This browse session request may include one or more addresses or references to various network resources or other content requested by the client computing device **102**. In an illustrative embodiment, the browse session request is transmitted in accordance with an API.

At block **906** the network computing provider **107** may select an associated NCC POP to instantiate a new browse session based on the browse session request. As discussed above with reference to FIG. **1**, a network computing provider **107** may include any number of NCC POPs distributed across any number of physical or logical locations. A network computing provider **107** may select a NCC POP to service a browse session request based on any number of factors, including, but not limited to available NCC POP resources (e.g., available memory, processor load, network load, etc.), a financial cost of servicing the browse session request at the NCC POP, the NCC POP location relative to a client computing device **102**, content provider **104**, or CDN POP **116**, a NCC POP cache status (e.g., whether a requested resource is already stored in an NCC POP cache), etc.

In one embodiment, the network computing provider **107** may select a number of NCC POPs to service a browse session request. For example, the network computing provider **107** may select two NCC POPs with different logical locations in the network. Each NCC POP may independently request and process network content on the behalf of the client computing device **102**, and the client computing device **102** may accept data from the first NCC POP to return a processing result. Subsequent to being selected by the network computing provider **107**, NCC POP **142** may obtain the browse session request. In one embodiment, NCC POP **142** may have the browse session request forwarded to it by a component of network computing provider **107**. In another embodiment, NCC POP **142** or client computing device **102** may receive connection information allowing the establishment of direct communication between NCC POP **142** and client computing device **102**. Illustratively, NCC POP **142** may be provided with the browse session request originally provided to network computing provider **107**, may be provided with a subset of information (e.g., just a network address of requested content), or may be provided additional information not included in the original browse session request.

Subsequent to the NCC POP **142** being selected, the network computing provider **107** may cause the NCC POP **142** to instantiate a new browse session. Illustratively, instantiat-

ing a new browse session instance may include loading a new virtual machine instance and/or browser instance at the NCC POP 142, reserving or allocating device memory, storage or cache space, processor time, network bandwidth, or other computational or network resources for the new browse session. Illustratively, one or more characteristics of the new browse session instance and/or browser instance may be based on client computing device 102 information included in the browse session request. For example, the browse session request may include a device type or browser type, a device screen resolution, a browser display area, or other information defining the display preferences or capabilities of the client computing device 102 or browser. The NCC POP 142 may accordingly instantiate a virtual machine instance and/or a browser instance with the same or similar capabilities as the client computing device 102. Illustratively, maintaining a virtual machine instance and/or browser with the same or similar capabilities as the client computing device 102 may allow the NCC POP 142 to process network content according to the appropriate dimensions and layout for display on the particular client computing device 102.

In some embodiments, the NCC POP 142 may utilize an existing virtual machine instance and/or browser instance in addition to, or as an alternative to, instating a new browse session. For example, subsequent to the NCC POP 142 being selected, the network computing provider 107 may cause the NCC POP 142 to associate an existing browser instance and/or virtual machine instance, such as one or more instances previously instantiated at the NCC POP 142, with the new browse session request. Illustratively, an existing browser session and/or virtual machine instance may correspond to another browse session, remote application session, or other remote process associated with the user or client computing device 102, or may be a previously instantiated software instance from an unrelated browse session or remote process. In other embodiments, the NCC POP 142 may instantiate a new browser or other application process in an existing virtual machine instance, or may combine the utilization of previously instantiated and newly instantiated software processes in any number of other ways. In still further embodiments, the network computing provider or NCC POP 142 may instantiate any number of new virtual machine instances and/or browser instances (or make use of existing instantiated instances) based on a single browse session request.

At block 908 the network computing provider 107 may provide a request for one or more network resources to a content provider or CDN service provider based on a network address included in the browse session request. In various embodiments, one or more network resources may be additionally or alternately retrieved from a cache local to the NCC POP 142 or otherwise associated with the network computing provider 107. One of skill in the art will appreciate that, in the case of other embodiments, the link or network address may correspond to a document or file stored in a digital file locker or other network storage location or at a cache component associated with the network computing provider 107 or client computing device 102. In some embodiments, the new session request may include a document or file in addition to or as an alternative to a network address. At block 910, the network computing provider 107 obtains the one or more network resources. Subsequent to obtaining the requested network resource, the NCC POP 142 may process the network resource to extract embedded resource identifiers.

At block 912, the network computing provider 107 may provide resource requests to one or more sources of content such as content providers, CDN service providers, and caches. The resource requests may correspond to embedded

resources based on the one or more embedded resource identifiers extracted from the one or more network resource as described in block 910 above. At block 914, the network computing provider 107 may obtain these embedded resources from any number of different sources, sequentially or in parallel.

At block 916, the network computing provider 107 may process the one or more network resources and associated embedded resources to determine a remote session browsing configuration for the processing and communication of content to the client computing device 102. A remote session browsing configuration may include any proprietary or public remote protocol allowing exchange of data and user interactions or requests between a client and a remote server. The remote session browsing configuration may illustratively include both a remote session communication protocol and a processing schema for providing processed (or unprocessed) content to a client computing device for display in the content display area of a browser.

Illustratively, a remote session browsing configuration may define or specify a remote session communication protocol, including, but not limited to, a network protocol, signaling model, transport mechanism, or encapsulation format for the exchange of state data, user interactions, and other data and content between the network computing provider and the client computing device. Examples of remote session communication protocols known in the art include Remote Desktop Protocol (RDP), X-Windows protocol, Virtual Network Computing (VNC) protocol, Remote Frame Buffer protocol, HTML, etc. For example, RDP illustratively specifies a number of processing mechanisms for encoding client input (e.g., mouse movement, keyboard input, etc.) into protocol data units for provision to a remote computing device, and corresponding mechanisms for sending bitmap updates and low level interface information back to the client device. As another example, the HTML protocol illustratively provides a mechanism for providing files defining interface information and containing resources references from a server to a client, and a corresponding mechanism for a client computing device to provide requests for additional files and resources to the server. In one embodiment, the NCC POP 142 may provide an initial communication to the client computing device 102 after determining the remote session communication protocol. This initial communication may allow the client computing device 102 to prepare to receive communications in the selected remote session communication protocol, and, in the case of pull remote session communication protocols like HTTP, may cause the client computing device to send an initial resource request to the browse session instance running on the NCC POP 142.

Each remote session browsing configuration may additionally define a split of processing actions between the network computing and storage service (e.g., NCC POP 142) and the client computing device (e.g., client computing device 102). In one embodiment, a particular split of processing actions may be based on or mandated by a particular remote session communication protocol. In another embodiment, a remote session communication protocol may allow several different splits of processing actions depending on the implementation or configuration of the protocol. For the purpose of illustration, many pieces of network content (e.g., Web pages, video, Flash documents) may require various processing actions before being displayed on a computing device. A Web page, for example, may be parsed to process various HTML layout information and references to associated resources or embedded content such as CSS style sheets and Javascript, as well as embedded content objects such as images, video, audio, etc.

The HTML and each referenced object or piece of code will typically be parsed and processed before a representative object model corresponding to the Web page may be constructed. This object model may then be processed further for layout and display in a content display area of a browser at the client computing device **102**. Illustrative browser processing actions are described in greater detail below with reference to FIG. **8**. One of skill in the art will appreciate that, in the case of other embodiments or applications, various other processing actions may be required.

A remote session browsing configuration may specify that various of the processing actions required for display of piece of network content be performed at the remote computing device, such as the NCC POP **142**, rather than at the client computing device **102**. Network content partially (or wholly) processed at the network computing provider may be referred to as a processing result. As discussed below, the split of processing actions may be associated with or linked to the remote session communication protocol used for exchanging data and client input between the NCC POP **142** and client computing device **102**.

For example, a remote session communication protocol such as RDP that transmits a processing result including low level interface information and bitmaps to the client computing device **142** for display may be associated with a remote session browsing configuration that specifies performing all, or nearly all, of the necessary content processing actions at the NCC POP **142**. While using RDP, the NCC POP **142** may, for example, run a full instance of a browser the NCC POP **142** and transmit a processing result consisting of bitmap updates corresponding to a representation of the displayed content to the client computing device **102**. The client computing device **102**, in this example, may merely be required to assemble the transmitted bitmap updates for display in the content display area of the browser, and may perform none of the processing of the actual HTML, Javascript, or data objects involved in the display of an illustrative piece of network content. As another example, a remote session browsing configuration utilizing a remote session communication protocol such as HTML may transmit network content in a largely unprocessed form. The client computing device **102** may thus perform all of the processing actions required for display of network content while the NCC POP **142** performs little or no processing.

The NCC POP **142** may base its determination of a remote session browsing configuration on any number of factors, including, but not limited to, one or more characteristics of one or more of the requested resources, content provider **104**, or CDN service provider **106**, one or more characteristics of the content address or domain, one or more characteristics of the client computing device **102**, browser or application, user, one or more characteristics of the NCC POP **142**, or one or more characteristics of the network or network connection, etc. Characteristics of requested resources may include, but are not limited to, a data format, a content type, a size, processing requirements, resource latency requirements, a number or type of interactive elements, a security risk, an associated user preference, a network address, a network domain, an associated content provider, etc. Characteristics of a content provider **104**, CDN service provider **106**, computing device **102**, or NCC POP **142** may include, but are not limited to, processing power, memory, storage, network connectivity (e.g., available bandwidth or latency), a physical or logical location, predicted stability or risk of failure, a software or hardware profile, available resources (e.g., available memory or processing, or the number of concurrently open software applications), etc. The NCC POP **142** may further consider

perceived security threats or risks associated with a piece of content or domain, preferences of a client computing device or a content provider, computing or network resource costs (e.g., a financial cost of processing or bandwidth, resource usage, etc.), predetermined preferences or selection information, any additional processing overhead required by a particular remote session browsing configuration, a cache status (e.g., whether a particular resources is cached at a NCC POP **142**, at the client computing device **102**, or at other network storage associated with the network computing provider), a predicted delay or time required to retrieve requested network content, a preferred content provider or agreements with a content provider for a particular remote session browsing configuration or level of service, a remote session browsing configuration being used for another (or the current) browse session by the same user, or any other factor.

In some embodiments, an NCC POP **142** may base a determination of a remote session browsing configuration on past behavior or practice. For example, an NCC POP **142** that has determined a remote browse session configuration for a particular resource in the past may automatically select the same remote browse session configuration when the resource is requested by the same (or potentially a different) user. As another example, a user that has a history of frequently accessing Web sites with extensive processing requirements may automatically be assigned a remote session browsing configuration that performs the majority of processing at the NCC POP **142**. In other embodiments, an NCC POP **142** may base a determination of a remote browse session configuration on predictions of future behavior. For example, an NCC POP **142** may base its determination of a remote browse session configuration for a particular resource on an analysis of past determinations made for a particular Web site, network domain, or set of related resources. A content provider that historically has provided video-heavy Web pages may be associated with a remote session browsing configuration that emphasizes video performance at the client computing device **102**. Illustratively, past historical analysis and future predictions may be considered as one or more of a number of factors on which to base the remote session browsing configuration determination process, or may be definitive in the decision making process. For example, once an NCC POP **142** determines a remote session browsing configuration for a particular content provider, it may skip the remote session browsing configuration determination process for any future resources served from the content provider. Illustratively, the NCC POP **142** may re-determine a remote session browsing configuration to be associated with the content provider after a fixed period of time, or after the NCC POP **142** has identified or determined a change in the content being served by the content provider.

In other embodiments, a network resource, Web site, network domain, content provider, or other network entity may specify or otherwise request the use of a particular remote browse session configuration in a resource tag, metadata, or other communication with an NCC POP **142**. The NCC POP **142** may treat the request as definitive, or may consider the request as one of multiple factors to be considered in the decision making process.

For example, a remote session browsing configuration utilizing a remote session communication protocol such as RDP may specify extensive processing to occur at the network computing provider **107** (e.g., at NCC POP **142**) rather than at the client computing device **102**. The remote session browsing configuration may thus leverage the processing power of the NCC POP **142** to achieve lower latencies and presentation delay when dealing with network content that requires a great

21

deal of pre-processing (e.g., content with a great deal of CSS or Javascript information defining page layout). The NCC POP 142 may therefore select a remote session browsing configuration that performs a substantial amount of processing at the network computing provider 107 and utilizes RDP or a similar remote session communication protocol for communication of processing-intensive content. Conversely, a remote session browsing configuration that utilizes a remote session communication protocol such as HTML may specify extensive processing at the client computing device 102 rather than at the network computing provider 107. The remote session communication protocol may thus achieve smaller delays and smoother presentation when presented with simple network content that requires very little processing or network content that requires rapid change in displayed content after its initial load. For example, a Web page with embedded video may perform better performing the majority of processing locally and utilizing HTML rather than RDP as a remote session communication protocol. A remote session browsing configuration specifying extensive processing at the network computing provider 107 must process the video at the NCC POP 142 and rapidly send screen updates (e.g. by RDP) to the client computing device 102, potentially requiring a great deal of bandwidth and causing choppy playback in the browser, while a remote session browsing configuration specifying local processing may provide raw video information directly to the client computing device 102 for display (e.g. by HTML), allowing for client side caching and a smoother playback of content.

As a further example, the NCC POP 142 in communication with a client computing device 102 with extremely limited processing power may elect to use a remote session browsing configuration that requires very little processing by the client computing device, for example, using RDP to transmit NCC POP 142 processed results. Conversely, an NCC POP 142 providing an extremely interactive Web page may elect to use a remote session browsing configuration that allows the client computing device 102 to handle user interactions locally in order to preserve interface responsiveness, for example, using HTML to transmit substantially unprocessed data. As a still further example, a NCC POP 142 may base the determination of a remote session browse configuration on preferences provided by the client computing device 102. A client computing device 102 may illustratively include preferences for a remote session browse configuration in an initial browse session request, or at any other time. The NCC POP 142 may utilize these preferences as an alternative to, or in addition to any other factor or decision metric. Illustratively, allowing the client computing device 102 to set or influence the selection of a remote session browse configuration allows the NCC POP 142 to take user preferences in account when determining a remote session browse configuration. For example, a user worried about initial page load times may prefer to use a remote session browsing configuration heavy on remote processing and utilizing an RDP remote session communications protocol, while a user wishing to maintain an extremely responsive interface may prefer using a remote session browsing configuration that performs the majority of the processing on the client computing device 102, for example, using an HTML remote session communication protocol.

Illustratively, the NCC POP 142 may base a determination of a remote browsing configuration on any factor or combination of factors. For example, the NCC POP 142 may select a remote session browsing configuration based on a single factor, or may assign weights to one or more factors in making a determination. In some embodiments, the determination process of the NCC POP 142 may change based on one or

22

more factors described above. For example, an NCC POP 142 communicating with a client computing device 102 over a network with a surplus of unused bandwidth may give a low weight to factors such as the network requirements of a remote browse session, and may give a higher weight to factors such as the latency of page interactions, while an NCC POP 142 communicating with a client computing device 102 over a limited bandwidth network may give a higher weight to factors dealing with the efficiency of the remote session browse protocol over a network.

In one embodiment, the NCC POP 142 may select a single remote session browsing configuration for a set of network content. For example, the NCC POP 142 may select a single remote session browsing configuration for a requested network resource such as a Web page. The NCC POP 142 may thus process the Web page together with all embedded content based on the selected remote browsing session protocol, and utilize the remote browsing session protocol to exchange user interaction data and updated browse session data for all embedded content associated with the Web page. In another embodiment, the NCC POP 142 may select different remote session browsing configurations for one or more resources in a set of network content. For example, a network resource such as a Web page may reference processing intensive embedded Javascript or CSS resources, as well as embedded video resources. The NCC POP 142 may select a first remote session browsing configuration for the Web page and all embedded resources excluding the embedded video resource, and a second remote session browsing configuration for the embedded video resource. Illustratively, this may result in the NCC POP 142 utilizing RDP to send a processing result to the client computing device 102 for display of the Web page and associated embedded resources, while utilizing HTTP to send the embedded video as a separate, unprocessed file. In one embodiment, the client computing device 102 may perform the minimal processing required to display the RDP processing result corresponding to the Web page and embedded resources, and may also perform additional processing necessary to display the embedded video, for example, overlaying the video on top of the displayed RDP representation of the Web page. Any number of remote session browsing configurations may be selected to correspond to any number of resources or objects included in a set of network content, regardless of whether resources or objects are obtained from a content provider 104 or CDN service provider 106 in one or more logical files or data structures.

Although the selection of a remote session browsing configuration is illustratively depicted herein as occurring after all network resources and associated embedded content have been obtained by the NCC POP 142, one skilled in the relevant art will appreciate that the selection of a remote session browsing configuration may be performed at any time. For example, the NCC POP 142 may select a remote session browsing configuration after receiving a new browse session request or related information from the client computing device, may select a remote session browsing configuration after obtaining a network resource, but before obtaining any associated embedded resources, or at any other time. In some embodiments, the NCC POP 142 may switch to a new remote session browsing configuration at some time subsequent to the client computing device 102 obtaining an initial processing result. Illustratively, the NCC POP 142 selecting a new remote session browsing configuration may occur automatically after a certain time period or event or in response to a change in network conditions, NCC POP 142 or client computing device 102 load or computing resources, or any other factor described above as potentially influencing the choice of

23

remote session browsing configuration. Illustratively, an NCC POP 142 dealing with other types or formats of information may select a remote session protocol based on any number of similar factors. For example, one of skill in the relevant art will appreciate that a similar schema may be developed for the processing of images, video, audio, data-
base information, 3 d design data, or any other file format or type of data known in the art.

The client computing device 102 may, in various embodiments, further instantiate a parallel browsing process sequentially or simultaneously with the request for a remote browse session. In one embodiment, a client computing device 102 may instantiate a traditional local browse session as known in the art (e.g., providing content requests from the browser and processing obtained resources locally) in addition to one or more remote browse instance executing at an NCC POP 142. In another embodiment, a client computing device 102 may be provided with unprocessed network resources by the NCC POP 142. Illustratively, the network resources may have been retrieved from one or more content providers, CDNs, or cache components by the NCC POP 142. The resources may be provided to the client computing device 102 to process locally in parallel with the remote browse instance executing at the NCC POP 142. In still further embodiments, the network computing provider or NCC POP 142 may instantiate any number of new virtual machine instances and/or browser instances (or make use of existing instantiated instances) to process resources and/or send processing results to the client computing device 102 in parallel. Illustratively, the local browse session at the client computing device 102 and the remote browse session instance at the NCC POP 142 may execute in parallel.

In one embodiment, a local browse session executing at the client computing device 102 may obtain unprocessed content (e.g., html Web pages, embedded content, and other network resources) from the NCC POP 142 responsive to a browse session request. Illustratively, the content may have been retrieved by the NCC POP 142 from a content provider, CDN, or cache in response to the browse session request. The unprocessed content provided by the NCC POP 142 may include all the content associated with the browse session request or may supplement content existing in a cache of the client computing device, retrieved from a content provider or CDN, or obtained from some other source. In one embodiment, a client computing device 102 may obtain all requested content from a local cache, and may not obtain any unprocessed resources or content from the NCC POP 142. Subsequent to obtaining the unprocessed content, client computing device 102 may process the requested content in parallel with a remote browse session executing at the NCC POP 142. For example, as the local browse session executing at the client computing device 102 is processing the requested content, a remote browse session executing at the NCC POP 142 may be processing the same content at substantially the same time. Once the NCC POP 142 has performed a set of processing actions on the content to generate a processing result (e.g., as specified by a determined remote session browsing configuration), the NCC POP 142 may provide the processing result to the client computing device 102.

For the purpose of illustration, a client computing device 102 may require a longer load time to obtain and process requested network resources than a browse session instance running at the NCC POP 142. For example, the NCC POP 142 may obtain and process content quickly due to its position on the network and the relative processing power of the local client computing device as compared to the NCC POP 142. Even if the NCC POP 142 provides the client computing

24

device 102 with all requested network content, the client computing device 102 may still obtain a processing result from NCC POP 142 before the local browse session has fully completed processing the requested resources. The client computing device 102 may complete any further processing steps and display the obtained processing result before completing local processing and display of the content. Illustratively, this may allow the client computing device 102 to take advantage of an NCC POP 142's quicker content load time relative to a traditional local browse session. Prior to the local browse session completing the processing all requested resources, the browser may process any user interactions locally and/or remotely as described in FIGS. 5A and 11 below.

Once the local browse session has fully obtained and processed resources corresponding to the requested content, the computing device 102 may determine whether to continue to display results obtained from the NCC POP 142 (and process user interactions at the NCC POP 142) using the determined remote session browsing configuration or switch to processing user interactions locally. Switching to process user interactions locally may include replacing a displayed representation of the requested resources based on a processing result obtained from the NCC POP 142 with a local display of the requested resources. For example, a browser may display a representation of a Web page corresponding to a processing result from the NCC POP 142 (e.g., RDP display information representing the rendered page) until the browser is finished processing and rendering the Web page locally. The browser may then replace the representation from the NCC POP 142 with the locally rendered representation of the Web page. Illustratively, replacing one representation with another representation may be transparent to the user. For example, the local and NCC POP 142 representations of the Web page may be identical or substantially identical. In one embodiment, when the NCC POP 142 representation of the web page is displayed, the browser may send various user interactions with the displayed page to the NCC POP 142 for processing. When the locally rendered version of the Web page is displayed, user interactions may be processed locally at the browser. Illustratively, the determination of which representation of the requested resources to display (e.g., local or from the NCC POP 142) may be based on any of the same factors described with reference to determining a remote session browse protocol in above.

In one embodiment, the client computing device 102 may switch to processing user interactions locally as soon as local resources are fully loaded. Illustratively, the remote browse session instance running at the NCC POP 142 may be terminated after switching to local processing, or the remote browse session instance may be maintained as a backup in case of unresponsiveness or a failure with regards to the local browse session. For example, the client computing device 102 may process user interactions locally, as well as sending remote user interaction data to the NCC POP 142 in accordance with the selected remote session browsing configuration. The remote user interaction data may be used by the NCC POP 142 to keep the remote browse session instance fully in parallel with the local browse process being executed by the browser at the client computing device 102. As long as the local browse session continues to handle user interactions, the NCC POP 142 may either refrain from sending updated processing results, or may send updated processing results ignored by the client computing device 102. If a problem develops with the local browse session at the client computing device 102, updated processing results may be provided to the client computing device 102 from the NCC POP 142 for

25

processing and display in lieu of the local browse session. Illustratively, this switch from the local browse session to remote processing may be transparent to the user. In some embodiments, the client computing device **102** may switch from a local browse session to a remote browse session instance based on factors other than unresponsiveness or failure at the local browser. For example, the client computing device **102** or network computing and storage component **107** may select between a remote and local browse session based on any of the factors enumerated with regards to determining a remote session browse protocol above

In another embodiment, the client computing device **102** may continue to process and display updated processing results from the NCC POP **142** even after the local browse session has fully loaded the requested content. The client computing device **102** may terminate the local browse session or may run the local browse session in parallel as a backup process in the converse of the example provided above. It should be appreciated that although the local browse session is described here for the purpose of illustration as being slower to load than the remote browse session instance, in some embodiments the local browse session may load the content faster than the remote browsing session, in which case the browser may process user interactions locally until the remote browse process has fully loaded the requested content. In some embodiments, the client computing device **102** may display and process user interactions through whichever browse session, local or remote, loads the requested content first.

In various other embodiments, the network computing provider **107** may instantiate multiple remote browse session instances to run in parallel in addition to or as an alternative to instantiating a local browse session. Illustratively, these parallel browse session instances may utilize any of the same or different remote session browse protocols, and may act as backups in the manner described above with regard to a local browse session, or may be used and switched between as alternatives in order to maximize browser performance at the client computing device **102**. For example, in response to one or more browse session requests, the network computing provider **107** may instantiate a browse session instance running on a first NCC POP and utilizing an RDP protocol as well as browse session instance running on a second NCC POP utilizing an X-Windows protocol. The client computing device **102** or the network computing provider **107** may determine which browse session instance and protocol should be used based on performance or resource usage considerations as described with regards to determining a remote session browse protocol above.

With continued reference to FIG. 9, at block **918**, the network computing provider **107** may process the obtained content, including the one or more requested network resources and embedded network resources, according to the determined remote session browsing configuration to generate an initial processing result. At block **920**, the network computing provider **107** may provide the initial processing result to the client for further processing and display in the content display area of the browser. For the purposes of further example, an illustrative client new browse session interaction routine **1000** implemented by client computing device **102** is described below with reference to FIG. 10. At block **922**, the start new browse session routine **900** ends.

FIG. 10 is a flow diagram illustrative of a client new browse session interaction routine **1000** implemented by client computing device **102**. New browse session interaction routine **1000** begins at block **1002** in response to an event or user request causing the client computing device **102** to load a

26

browser for viewing network content. At block **1004**, the client computing device loads locally managed components of the browser, including all local interface components. As described above with reference to FIGS. 5A and 7, local interface components may include toolbars, menus, buttons, or other user interface controls managed and controlled by the software browser application or any other process executing or implemented locally at the client computing device. At block **1006**, the client computing device **102** provides a request for a new browse session instance to the network computing provider **107**. From the perspective of the user of the client computing device, the request for the new browse session corresponds to the intended request to transmit the request to one or more corresponding content providers **104**. In other embodiment, the new session request may correspond to a request to load a file or other document (e.g., a request to load an image in a photo-editing application, etc.). Illustratively, the request may be generated automatically as a result of the browser loading (e.g., a request for a default or “home” page), or may be generated as a result of a user following a link or entering a network address into an address bar. As illustrated with respect to FIG. 2, the browse session request is transmitted first to a network computing provider **107**. In an illustrative embodiment, the network computing provider **107** utilizes a registration API to accept browse session requests from the client computing device **102**.

A browse session request may include any number of pieces of data or information including, but not limited to, information associated with a user, information associated with the client computing device **102** or software on the client computing device (e.g., hardware or software information, a device physical or logical location, etc.), information associated with the network **108**, user or browser preferences (e.g., a requested remote session browse protocol, a preference list, a decision tree, or other information), information associated with the network computing provider **107**, information associated with one or more pieces of requested network content (e.g., the network address of a network resource), etc. For example, a browse session request from the client computing device **102** may include information identifying a particular client computing device hardware specification or a hardware performance level, latency and bandwidth data associated with recent content requests, a desired security level for processing different types of content, a predetermined preference list of remote session browse protocols, and one or more network addresses corresponding to requested network resources, among others. In another example, the browse session request can include information identifying a client computing device **102** screen resolution, aspect ratio, or browser display area in the browse session request may allow the network computing provider **107** to customize the processing of network content for display on the client computing device. As previously described, the browse session request can include network address information corresponding to a requested network resource, which may be in any form including, but not limited to, an Internet Protocol (“IP”) address, a URL, a Media Access Control (“MAC”) address, etc. In one embodiment, the request for a new browse session instance may correspond to the network computing provider receiving a request for a new browse session instance at block **904** of FIG. 9 above.

At block **1008**, the client computing device **102** obtains an initial processing result from the network computing provider **107**. Illustratively, the format and data included in the initial processing result may vary based on the remote session browsing configuration selected by the network computing provider **107**. In one embodiment, the initial processing result

may include or be preceded by data informing the client computing device **102** of the choice of remote session browsing configuration and/or establishing a connection over the remote session communication protocol corresponding to the selected remote session browsing configuration. As discussed above with reference to FIGS. **8** and **9**, the obtained initial processing result may include requested content with one or more processing actions performed by the network computing provider **107**. Subsequent to obtaining the initial processing result, the client computing device **102** may perform any remaining processing actions on the initial processing result at block **1010**.

At block **1012**, the client computing device **102** displays the content corresponding to the processed initial processing result. For example, the client computing device **102** may display the processed client in the content display area **702** of a browser **700** as described in FIG. **7A** above. In one embodiment, the processing result may only include display data corresponding to content displayed by a browser, and may not include display data corresponding to, for example, the interface controls of a browser instance at the NCC POP **142**, the desktop of a virtual machine instance corresponding to the browse session, or any other user interface of the NCC POP **142**. For example, the NCC POP **142** may process a Web page and associated content for display via RDP in a browser instance running in a virtual machine instance at the NCC POP **142**. The browser instance may have one or more interface elements such as toolbars, menus, scroll bars, etc., in addition to the displayed Web page. The NCC POP **142** may send an RDP processing result corresponding to the displayed Web page only, without any of the interface elements associated with the browser. Illustratively, including an RDP processing result corresponding to the displayed Web page only may allow the browser at the client computing instance **102** to display the Web page by assembling the RDP processing result in the content display area of the browser without any further processing. In another embodiment, the RDP processing result may include a full virtual machine desktop and browser window corresponding to the full interface displayed at the NCC POP **142** browse session instance. The client computing device may automatically identify the area of the RDP processing result corresponding to the requested content, and may display only this area in the content display area of the browser.

At block **1014**, the client computing device **102** processes local and remote user interactions. An illustrative routine for processing user interactions is provided below with reference to FIG. **11**. At block **1016** the routine ends. Illustratively, a browse session instance instantiated by the network computing content provider **107** may terminate when a browser window or content display area is closed, may terminate when a remote session browse protocol is replaced by a parallel process at the client computing device **102**, or may terminate in accordance with a timer or other event. Illustratively, if a browse session has terminated automatically due to a time-out but has associated content still displayed in a browser at the client computing device **102**, later attempts by the user to interact with the content may result in a new browse session request being provided to the network computing and storage service provider **107** to start a new browse session according to the last state of the terminated session. Illustratively, terminating a remote browse session after a time-out may allow the network computing storage provider **107** to save computing resources at the NCC POP. In one embodiment, this process may be transparent to the user at client computing device **102**, even though the remote browse session has been terminated during the intervening period.

FIG. **11** is a flow diagram illustrative of a process user interaction routine **1100** implemented by a client computing device **102**. Process user interaction routine **1100** begins at block **1102** in response to an interaction by a user. Illustratively, process user interaction routine **1100** may begin subsequent to the display of content in a content display area of a browser interface. For example, process user interaction routine **1100** may correspond to block **1014** of FIG. **10** above.

Illustratively, the displayed content may have one or more interactive elements, such as forms, buttons, animations, etc. User interaction with these interactive elements may require processing and display of updated content in the content display area. For example, selecting an element in a drop-down menu on a Web page may require processing and may change the configuration or visual appearance of the Web page or embedded resources. Illustratively, the processing required by user interaction with the displayed content may be handled as a local user interaction at the client computing device **102** or as a remote user interaction at the NCC POP **142** depending on the remote session browsing configuration in use. For example, if a remote session browsing configuration utilizing substantial local processing (e.g., sending unprocessed files over HTML), user interactions with displayed content may typically be handled as local user interactions at the client computing device **102**. Illustratively, handling user interactions with displayed content as local user interactions at the client computing device **102** may allow for better responsiveness and fewer delays with simple user interactions (e.g., selection of a radio button, or typing text into a field), as interaction data corresponding to the interaction does not need to be sent to the NCC POP **142** for processing.

As a further example, if a remote session browsing configuration utilizing heavy remote processing of content (e.g., sending processed bitmap data over RDP) is being used as the remote session browsing configuration, all user interactions with displayed content may be handled as remote user interactions. For example, user input (e.g., keyboard inputs and cursor positions) may be encapsulated in RDP protocol data units and transmitted across network **108** to the NCC POP **142** for processing. Illustratively, the NCC POP **142** may apply the user interactions to the network content and transmit processing results consisting of updated bitmaps and interface data corresponding to an updated representation of the content back to the client computing device **102**. Illustratively, handling user interactions with displayed content as remote user interactions at the NCC POP **142** may have a negative impact on interface responsiveness, as data is required to pass over the network and is limited by network latency; however, user interactions that require a substantial amount of processing may perform better when handled as remote user interactions, as the processing latency of the NCC POP **142** may be substantially lower than the processing latency of the client computing device **102**.

In addition to a content display area for displaying network content, a browser may have one or more local interface components, such as toolbars, menus, buttons, or other user interface controls. Interactions with local interface components may be treated as local user interactions or remote user interactions depending on the processing required by the interaction and the remote session browsing configuration as further depicted in illustrative FIG. **7A**. For example, some local interface components may be managed locally by browser code running on the client computing device, while other local interface components may have one or more locally managed aspects (e.g., button click feedback, scroll bar redraw, etc), and one or more remote managed aspects

treated as remote user interactions (e.g., page refresh, requesting a page at an address in an address bar, etc.)

At block **1104**, the client computing device **102** obtains a user interaction from the user. This user interaction may be an interaction with local interface components as described in FIG. 7A and above, or may be an interaction with any interactive elements of the content displayed in the content display area of the browser, such as form fields, buttons, animations, etc. User interaction with these local interface components or interactive elements of displayed content may require local and/or remote processing depending on the nature of the component or element and the processing split specified by the remote session browsing configuration as described in FIG. 7A and above. At block **1106**, the client computing device **102** determines the interaction processing requirements for the obtained user interaction. At decision block **1108**, if the user interaction has local aspects (e.g., button click feedback, a change to a local browser state, a content element being processed at the client computing device, etc.) the routine **1102** moves to block **1110** to process the local aspect or aspects of the user interaction at the client computing device **102** and subsequently update the local interface components at block **1112**. Illustratively, and as discussed above, aspects of the interaction and updating interface components and elements locally allows a browser to provide responsive user interfaces and content. Subsequent to processing local aspect(s) of the user interaction, or if the user interaction has no local elements (e.g., a user interaction with a content element displayed in the content display area when using a remote session browsing configuration processing entirely on the server side and utilizing an RDP remote session communication protocol) the routine **1102** moves to decision block **1114**. If the user interaction has remote aspects that require processing, the routine **1102** moves to block **1116** and provides remote user interaction data to the network computing provider **107**. Illustratively, in the case of a heavily server side remote session browsing configuration utilizing an RDP remote session communication protocol, the remote user interaction data may include input data such as a cursor position or keyboard input encapsulated in one or more RDP protocol data units. In some embodiments of remote session browsing configurations utilizing RDP or other remote session communication protocols, particular aspects of remote user interaction data such as cursor positions may be provided to the network computing provider **107** on a continuous basis, while in other embodiments of remote session browse configurations remote user interaction data may only be provided to the network computing provider **107** when associated with a user interaction that requires remote processing.

At block **1118**, the client computing device **102** obtains an updated processing result from the network computing provider **107**, the network computing provider **107** having processed the remote user interaction data to generate an updated representation of the content. At block **1120**, the client computing device **102** performs any additional processing required on the updated processing result (based on the remote session browsing configuration) and at block **1122** displays the updated processing result in the content display area of the browser. At block **1124** the process user interaction routine **1102** ends. Illustratively, the routine may be executed again any number of times in response to further user interactions with the browser and displayed content.

FIGS. 12 and 13 illustrate routines **1200**, **1300** implemented by a network computing provider **107** for receiving and processing narrations of network resources, and for providing those narrations to end users. Network resources, such as web pages, often contain text, images, videos, animations,

applets, input fields, and the like. Navigating and interacting with such web pages can be difficult for users with visual and other impairments. For example, some users of client computing devices **102** may not be able to read the text on a web page due to poor eyesight, inability to read, limited display capabilities, a need to focus visual attention elsewhere, etc. Spoken narration can be provided to such users, and can include not only a verbatim reading of any textual components of the web page, but also a description of the images, animations, and other non-textual visual elements. Additionally, narrated previews can be provided to the user, describing what will happen if a user activates a link or other input field, etc. In some embodiments, a narrator can paraphrase the text in a web site or provide analysis and opinion about elements of the web page to the user in addition to, or instead of, a strictly factual recitation of the visual elements of the web page. In some embodiments, the spoken narration can be generated by text-to-speech software. In some embodiments, the narration is not a spoken description of visual elements, but rather a written or otherwise visual description. Such a visual narration may be provided to users with hearing impairments, to users who required additional help or instruction on how to interact with elements on a page, to users who wish to receive the commentary and analysis described above without requiring use of speakers or headphones, etc. In some embodiments, a visual narration can be coupled to a spoken narration to provide a rich, multimedia narrated experience to a user.

Narrations can be provided by users of a client computing device **102**, utilizing a browser interface as illustrated and described above with respect to FIG. 7A. In some embodiments, narrations can be provided using alternative techniques, such as stand-alone narration software, a remote desktop connection into a network computing provider **107**, etc. The discussion that follows with respect to FIG. 12, describing spoken narrations with visual elements recorded on a client computing device **102** and transmitted upon completion to a network computing provider **107**, is illustrative only, and is not intended to be limiting.

FIG. 12 illustrates a routine **1200** implemented by a network computing provider **107** for receiving narrations from client computing devices **102**, and the storage of the narrations so that the narrations may be provided to users of client computing devices **102** at a later time. The routine **1200** begins at block **1202** with the initialization of a narration creation session. A narration can be created by a narrator, which may be a single person, a team of narration creation specialists, text-to speech software, some combination thereof, or any other appropriate narration creator. In the example described herein with respect to FIG. 12, the narrator can be a single user of a client computing device **102**. A narrator may decide to provide a narration of a network resource, such as a web page, for one of any number of reasons. For example, a narrator may provide narrations for pay or other compensation, on a volunteer basis, etc. The narrator may utilize a browser interface **700** as described above and illustrated in FIG. 7A. The client computing device **102** may connect to a network computing provider **107**, and in response the network computing provider **107** may instantiate a network browser application which communicates with the browser application executing on the client computing device **102**, as describe in detail herein. The initialization of a narration creation session may involve fewer or additional procedures not detailed here. In response to the successful initialization of a narration creation session, the routine **1200** can proceed to block **1204**.

31

At block **1204**, the narrator can select a network resource, such as a web page, for which to create a narration. For example, if a narrator wishes to create a narration for a news web site, the narrator may go to the home page of the news web site and transmit a communication to the network computing provider **107** that the narrator has selected the home page of the news web site. This communication can be transmitted as part of, or in addition to, the other communications between the network browser of the network computing provider **107** and the browser application on the client computing device **102**, described in detail above. For example, a browser application on the client computing device **102** transmits requests for network resources to the network computing provider **107**, which then retrieves, processes, and provides the network resources to the client computing device **102**, while in some cases retaining some or all of the processing of the network resource at the network computing provider **107**. In this example, a narrator may request the home page of the news web site, and the network computing provider **107** may retrieve, process, and transmit the home page to the client computing device **102**. The client computing device **102** may then transmit additional communications to the network computing provider **107** detailing user interactions with the home page, such as the activation of the create narration button **704** illustrated in FIG. 7A. The network computing provider **107** may then be notified that the user is creating a narration for the currently displayed web page. In response to receiving the notification, the routine **1200** can proceed to block **1206**.

At block **1206** of FIG. 12, the narrator may record and transmit to the network computing provider **107** a spoken overview of the currently displayed web page. Returning to the example above, the narrator may describe the overall visual layout of the page and read any top-level headlines. The page overview narration can be created as a separate recording, or can be part of a larger running recording of the narration of the entire page. The recording can be streamed to the network computing provider **1206** for storage as it is being recorded. In some embodiments, the recording can be made and saved as a file or buffer on the client computing device **102**, which is later transmitted to the network computing provider **107**. In some embodiments and in some specific cases, a page overview may be neither required nor appropriate, depending upon the implementation and the specific network resource being narrated. In response to finishing the page overview, or if none is provided, the routine **1200** can proceed to block **1208**.

At block **1208**, the narrator can select a specific portion of the web page to narrate. This can be useful, for example, when a page has multiple distinct and potentially unrelated portions. The web page illustrated in FIG. 7A, for example, includes a dynamic top stories portion **720** and a best seller list portion **708**, in addition to links, headlines, paragraphs of text, etc. Each portion can be given a separate narration, allowing a listening user to select and listen to only narrations of portions that the listening user is interested in. Separating the narration of the web site into portions can also help to ensure that if one portion of the web site is changed, only the corresponding portion narration will need to be removed or replaced, rather than all narrations of the web page. In response to a portion of the web page being selected, execution of the routine **1200** can proceed to block **1210**.

At block **1210**, the narrator can provide an overview of the portion, in a manner similar to the page overview discussed above and shown at block **1206**. The recording can be streamed to the network computing provider **1206** for storage as it is being recorded. In some embodiments, the recording can be made and saved as a file or buffer on the client com-

32

puting device **102**, which is later transmitted to the network computing provider **107**. Additionally, the portion overview may not be appropriate or required, depending on the implementation and the network resource. In response to the completion of the portion overview, or if none is provided, the routine **1200** can proceed to block **1212**.

At block **1212** the narrator can provide a detailed narration of the selected portion. For example, in the news web site example, the narrator may have selected the best seller list portion **708** illustrated in FIG. 7A. The narrator can read the contents of the list, and can also describe what will happen if the user were to click on an item in the list. For example, the best seller list portion **708** may contain images **710** which are clickable, and which take the user to a more detailed discussion of the list item, or to another page where the user can purchase the list item, etc. In some embodiments, the narrator can activate the link and provide the user with a narrated preview of the linked page. In response to completing the narration of a portion, the routine **1200** can proceed to decision block **1214**.

At decision block **1214** of FIG. 12, the narrator may choose to narrate additional portions, if any exist, of the active network resource. Returning to the news web site example, the narrator may wish to narrate the top news stories portion **720** illustrated in FIG. 7A. In such a case, the routine **1200** could return to block **1208**, where the narrator could select the top news stories portion **720** for narration. If the narrator does not wish to provide narration of additional portions, or if there are no more portions to narrate, the routine **1200** can proceed to decision block **1216**.

At decision block **1216** of FIG. 12, the narrator may chose to narrate additional web pages. For example, in the news web site example, the narrator may wish to narrate detail pages dedicated to a particular news story. In such a case, the routine **1200** could return to block **1204**, where the narrator could select the next web page to narrate. If the narrator does not wish to provide narration of additional pages, or if there are no more pages associated with the web site, the routine **1200** can proceed to block **1218**, where execution terminates.

FIG. 13 illustrates a routine **1300** implemented by a network computing provider **107** for providing narrations to users of client computing devices **102**. The routine **1300** begins at block **1302** with the receipt of a narration request from a client computing device **102**. For example, a user of a client computing device **102** may navigate to a network resource, such as the news web site described above, and activate a narration playback button **712** or link **724**, as illustrated in FIGS. 7B and 7C. In response to a page narration request, the routine **1300** can proceed to block **1304**.

At block **1304**, the network computing provider **107** can cause the page overview narration, if there is one, to play back on the client computing device **102**. The narration can be transmitted as a file to the client computing device **102** for playback, can be streamed to the client computing device **102**, etc. As described above, the narration can include both audio and visual elements, and the narration can be played back on the client computing device **102** without interfering with the user's ability to scroll, activate links, or otherwise interact with the page as normal. In some embodiments, the user may be provided with an option to download the narration rather than having the narration played back interactively. For example, an audio or video file may be provided so that the user can watch or listen to the narration offline.

At block **1306** the network computing provider **107** can determine which changes, if any, have occurred to the network resource since the narration was recorded. Returning to the news web site example, the network computing provider

33

107 can use one or more of the techniques 1308 to determine which changes have occurred to the top stories portion 720 of FIGS. 7A, 7B and 7C. As illustrated in FIG. 7A, the top stories portion 720 may have contained one or more content portions when the narration was created, and as illustrated in FIG. 7C, the top stories portion 720 may contain one or more different content portions when narration playback is requested. The network computing provider 107 can utilize techniques 1308 of FIG. 13, such as processing user notifications, performing pixel analysis, comparing digital page fingerprints, etc., or any combination thereof, to determine which portions, if any, have changed since the narration was created. The determination of changes can be made when the page narration playback request is initially received at block 1302, during page overview playback at block 1304, in response to receiving a request for playback of the narration of a specific section at decision block 1310, described below, or at any other appropriate time.

In one example, the network computing provider 107 may receive notifications from multiple users of client computing devices 102 when the users are playing back narrations. The notifications can inform the network computing provider 107 that one or more portions of a network resource have changed. The network computing provider 107 can perform some processing and analysis on the notifications, such as applying machine learning algorithms to the notifications and to the past and present pages, giving weight to notifications that are corroborated by subsequent notifications, assigning staff members to inspect the page to confirm that the notifications are accurate, etc. In another example, the network computing provider 107 can perform pixel sampling analysis on the current version of the web page, sampling pixels and determining what content is to be displayed on the screen at the pixel location by analyzing the network browser's rendering of the web page, receiving pixel data from the client computing device 102, etc. The content at the pixel can be determined by consulting the HTML tag that governs the display at the pixel location, by saving the actual value of the color displayed by the pixel, or some other technique. The network computing provider 107 can then perform the same pixel sampling analysis to the version of the web page as it existed when the narration was created. This can involve restoring a copy of the web page from the time the narration was created, loading pixel profile data that was stored at the time the narration was stored, etc. By comparing profile data from the current web site to the pixel profile data from the web page as it existed when the narration was created, the network computing provider 107 can determine differences between the two, and the difference may represent portions of the web page that have changed. In yet another example, the network computing provider may create or observe a digital web page fingerprint of the page at the time the narration was created, and can compare the fingerprint to a digital web page fingerprint of the page at the time playback of the narration is requested. The fingerprint can be a hash of the page or each portion of the page, a pixel/content map similar to the pixel analysis described above, meta data about the page, some combination thereof, etc.

At decision block 1310, the network computing provider 107 can receive a selection of a portion or specific element on the web page. If a selection is received, execution of the routine 1300 can proceed to block 1312. If no selection is received, execution of the routine 1300 can proceed to decision block 1316.

At block 1312, the network computing provider 107 can notify the user of any changes to the portion which have occurred since the time the narration was created. The noti-

34

fication can be an audio and/or visual notification, and in some case can include information about what has changed. For example, in the news web site illustrated in FIGS. 7A and 7C, the user can be notified that the content in the top stories portion 720 has been changed, and that the narration playback that follows may not be accurate. In some embodiments, rather than notifying the user that the portion has changed, the out of date portion of the narration can be withheld from playback, either with or without a notification to the user. In some embodiments, the network computing provider 107 can monitor changes to underlying network resources as a stand-alone server process without requiring a request from a client computing device 102, and can solicit an updated narration for out of date portions before they are requested by a user. Newly created narration portions can then be incorporated into the narration for the web page.

If a narration exists for the selected portion, the routine 1300 can proceed to block 1314, where the network computing provider 107 causes playback of the narration to the user of the client computing device 102. The narration can be played back in a manner similar to the page overview playback detailed above with respect to block 1302. As described above, the user of the client computing device 102 may be permitted to interact with the web page during narration playback, such as clicking on links, etc.

At decision block 1316, the network computing provider 107 can receive user notifications. As described above with respect to the techniques of block 1308, the network computing provider 107 may analyze user notifications transmitted to the network computing provider 107. A user notification may correspond to a portion of the web page to which the played narration no longer applies, is erroneous, etc. For example, a user of a client computing device 102 may determine that a portion of the web page being narrated, such as the top stories portion 720 illustrated in FIG. 7C, has changed since the narration was created, and the narration no longer applies to the portion. The user may then transmit a notification to the network computing provider 107. If a user notification is received, the routine 1300 can proceed to block 1318. Otherwise, execution of the routine can proceed to block 1320.

At block 1318, the network computing provider 107 can process a notification received in decision block 1316. Processing can include comparing the current version of the web page to an archived version of the page as it existed when the narration was originally created, dispatching a staff member to manually confirm any change in the page, etc. If a change is confirmed, the network computing provider 107 can alter the narration, as described above. The routine 1300 can then proceed to block 1320, where execution terminates.

FIG. 14 illustrates a routine 1400 implemented by a network computing provider 107 for updating a narration of a network resource in response to changes in the content of the network resource. The routine 1400 begins at block 1402 with the initialization of a narration edit session. The narration edit session may be initialized in response to one of any number of events. For example, a user may send a notification that a narration, or portion thereof, is no longer accurate, an automated process may monitor the content of narrated pages to determine when a change has been made, etc. Additionally, narration information may be associated with expiration criteria, such as time-based criteria or event-criteria, that would result in the determination that some portion of the narration is no longer valid. In response to the successful initialization of a narration edit session, the routine 1400 can proceed to block 1404.

35

At block 1404, the network computing provider 107 can retrieve the narrated web page and determine whether a portion of the page has changed since the narration was created or last updated. The network computing provider 107 can use a number of techniques to determine whether the page has changed, as described above with respect to FIG. 13. For example, the network computing provider 107 can perform pixel analysis, compare page fingerprints, etc. In response to determining that a portion of the page has changed, the routine 1400 can proceed to block 1406.

At block 1406, the network computing provider 107 can request an updated narration for the changed portion. In some embodiments, the network computing provider 107 may generate a request to the same narrator that created the original narration, or to any of a group of predetermined narrators. Alternatively, the network computing provider 107 can place a request in a queue, where the request can be fulfilled by the next available narrator. In some embodiments, the network computing provider 107 can utilize text-to-speech software to generate an updated aural narration.

At block 1408, the network computing provider 107 can receive an updated narration in response to the request generated in block 1406. The routine 1400 can then proceed to block 1410. In some embodiments, a substantial period of time may pass between generation of a request for an updated narration and receipt of the updated narration, or it may be critical to update the narration to remove the inaccurate portion even if an updated narration does not replace the inaccurate portion. In such cases, the routine 1400 may proceed to block 1410 without receiving the updated narration, and can return to block 1408 when the updated narration has been received.

At block 1410, the network computing provider 107 can modify the narration by adding the updated narration portion received in block 1408. For example, the network computing provider 107 can replace the out-of-date portion with the updated narration portion. In some embodiments, an updated narration for the portion may not be available, such as when a request for an updated narration has been generated in block 1406 but no narrator has created an updated narration. In such cases, the network computing provider 107 can remove the outdated portion of the narration without replacing, or by replacing it with a notification to the user that there is no narration available. In response to modifying the narration, the routine 1400 can proceed to decision block 1412.

At decision block 1412 the network computing provider 107 can determine whether there are additional portions of the web page for which the corresponding narration is no longer accurate. If so, the routine 1400 can return to block 1406 for each of the portions. If there are no additional portions of the narration known to be inaccurate, the routine 1400 can proceed to block 1414, where the updated narration can be finalized and execution of the routine 1400 terminates.

While illustrative embodiments have been disclosed and discussed, one skilled in the relevant art will appreciate that additional or alternative embodiments may be implemented within the spirit and scope of the present invention. For example, the techniques described herein may be utilized, without departing from the scope of the present invention, to allow remote processing management in any number of other software applications and processes, including, but not limited to, image or video editing software, database software, office productivity software, 3 d design software, audio and sound processing applications, etc. Additionally, although many embodiments have been indicated as illustrative, one skilled in the relevant art will appreciate that the illustrative embodiments do not need to be combined or implemented

36

together. As such, some illustrative embodiments do not need to be utilized or implemented in accordance with scope of variations to the present disclosure.

Conditional language, such as, among others, “can,” “could,” “might,” or “may,” unless specifically stated otherwise, or otherwise understood within the context as used, is generally intended to convey that certain embodiments include, while other embodiments do not include, certain features, elements and/or steps. Thus, such conditional language is not generally intended to imply that features, elements and/or steps are in any way required for one or more embodiments or that one or more embodiments necessarily include logic for deciding, with or without user input or prompting, whether these features, elements and/or steps are included or are to be performed in any particular embodiment.

The terms “comprising,” “including,” “having,” and the like are synonymous and are used inclusively, in an open-ended fashion, and do not exclude additional elements, features, acts, operations, and so forth. Also, the term “or” is used in its inclusive sense (and not in its exclusive sense) so that when used, for example, to connect a list of elements, the term “or” means one, some, or all of the elements in the list. Conjunctive language such as the phrase “at least one of X, Y and Z,” unless specifically stated otherwise, is otherwise understood with the context as used in general to convey that an item, term, etc. may be either X, Y or Z. Thus, such conjunctive language is not generally intended to imply that certain embodiments require at least one of X, at least one of Y and at least one of Z to each be present.

Any process descriptions, elements, or blocks in the flow diagrams described herein and/or depicted in the attached FIGURES should be understood as potentially representing modules, segments, or portions of code which include one or more executable instructions for implementing specific logical functions or steps in the process. Alternate implementations are included within the scope of the embodiments described herein in which elements or functions may be deleted, executed out of order from that shown or discussed, including substantially concurrently or in reverse order, depending on the functionality involved, as would be understood by those skilled in the art. It will further be appreciated that the data and/or components described above may be stored on a computer-readable medium and loaded into memory of the computing device using a drive mechanism associated with a computer readable storing the computer executable components such as a CD-ROM, DVD-ROM, or network interface further, the component and/or data can be included in a single device or distributed in any manner. Accordingly, general purpose computing devices may be configured to implement the processes, algorithms, and methodology of the present disclosure with the processing and/or execution of the various data and/or components described above.

It should be emphasized that many variations and modifications may be made to the above-described embodiments, the elements of which are to be understood as being among other acceptable examples. All such modifications and variations are intended to be included herein within the scope of this disclosure and protected by the following claims.

What is claimed is:

1. A system including at least one processor for providing narration of dynamic network content, the system comprising:
a network computing component executing on one or more physical network computing devices, wherein the one or

37

more physical network computing devices are programmed by executable instructions to at least: receive a web page narration from a first client computing device, the narration comprising a plurality of audio clips including a first audio clip, each of the plurality of audio clips comprising an aural description, by a user of the first client computing device, of a portion of the web page, wherein the web page narration is generated separately from the web page; determine whether a portion of the web page has changed after the narration comprising the first audio clip describing the portion was received, wherein the one or more physical network computing devices are programmed to determine whether the portion of the web page has changed by at least: receiving change information comprising at least one of: a user notification regarding a change to the portion of the web page; a pixel analysis of the portion of the web page, or a hash of the portion of the web page; and in response to receiving the change information, comparing a present version of the portion with a previously received version of the portion; in response to determining that the portion of the web page has changed: obtain a second audio clip describing an updated version of the portion of the web page; and replace the first audio clip with the second audio clip describing the updated version of the portion of the web page; cause playback, on a second client computing device, of at least a portion of the web page narration, the portion of the web page narration comprising at least the second audio clip, wherein the playback is caused in response to receiving a user interaction from the second client computing device; and enable a user of the second client computing device to interact directly with the web page during playback of the at least one of the plurality of audio clips.

2. The system of claim 1, wherein the aural description of the portion of the web page comprises at least one of a verbal transcription of text displayed in the portion, a verbal description of textual content displayed in the portion, or a verbal description of visual content displayed in the portion.

3. The system of claim 1, wherein the web page narration includes at least one visual indication of the portion of the web page described by at least one of the plurality of audio clips.

4. A computer-implemented method for providing narration of dynamic network content, the method comprising: under the control of a network computing component executing on one or more physical computing components of a network computing provider, the physical computing components configured to execute specific instructions, receiving a network resource narration from a first client computing device, the network resource narration corresponding to a first aural presentation of at least a portion of the network resource, wherein the network resource narration is generated separately from the network resource; determining whether a portion of the network resource has changed since the network resource narration comprising the first aural presentation of the portion was received, wherein the determining comprises: receiving an indication of at least one of: a user notification regarding the portion of the network

38

resource; a pixel analysis of the portion of the network resource; or a hash of the portion of the network resource; and in response to receiving the indication, comparing a present version of the portion with a previously received version of the portion; in response to determining that the portion of the network resource has changed: obtain a second aural presentation of an updated version of the portion of the network resource; and replace the first aural presentation with the second aural presentation of the updated version of the portion of the network resource; causing playback, on a second client computing device, of at least a portion of the network resource narration comprising the second aural presentation, wherein the playback is caused in response to receiving a user interaction from the second client computing device; and enabling a user of the second client computing device to perform actions available to users of the network resource when no network resource narration is played back.

5. The computer-implemented method of claim 4, wherein the network resource comprises at least one of a web page, image, video, applet, or document.

6. The computer-implemented method of claim 4, wherein the portion of the network resource narration comprises at least one of an aural description, an aural recitation, an aural translation, a textual transcription, a textual translation, a textual description, a visual transcription, or a visual annotation of a portion of the network resource.

7. The computer-implemented method of claim 6, wherein at least one visual transcription of the portion of the network resource comprises a sign language translation of at least a portion of a spoken element of the network resource.

8. The computer-implemented method of claim 6, wherein at least one aural description of the portion of the network resource comprises a verbal description of an image or video displayed in the portion of the network resource.

9. The computer-implemented method of claim 6, wherein at least one aural description of the portion of the network resource comprises a verbal description of a result of an action taken on an element displayed in the portion of the network resource.

10. The computer-implemented method of claim 9, wherein the element comprises one of a link, button, drop-down, menu, or input control, and wherein the action taken comprises one of clicking on, typing in, or activating the element.

11. The computer-implemented method of claim 4, wherein the network resource narration further comprises at least one visual effect.

12. The computer-implemented method of claim 11, wherein the visual effect indicates the portion of the network resource described by the first aural presentation.

13. A system including at least one processor for providing narration of interactive network content, the system comprising: a network computing component executing on one or more physical network computing devices, wherein the one or more physical network computing devices are programmed by executable instructions to at least: receive a narration from a first client computing device, the narration comprising a first audio clip describing a portion of the network resource, wherein the narration is generated separately from the network resource;

39

determine whether a portion of the network resource has changed since the narration comprising the first audio clip describing the portion was received, wherein the one or more physical network computing devices are programmed to determine whether the portion of the network resource has changed by at least:

receiving an indication of at least one of: a user notification regarding the portion of the network resource; a pixel analysis of the portion of the network resource; or a hash of the portion of the network resource; and

in response to receiving the indication, comparing a present version of the portion with a previously received version of the portion;

in response to determining that the portion of the network resource has changed:

obtain a second audio clip describing an updated version of the portion of the network resource; and replace the first audio clip with the second audio clip describing the updated version of the portion of the network resource; and

cause playback, on a second client computing device, of at least a portion of the narration comprising at least the second audio clip.

14. The system of claim 13, wherein the network resource comprises at least one of a web page, image, video, applet, or document.

15. The system of claim 13 wherein the portion of the narration comprises at least one of an aural description, an aural recitation, an aural translation, a textual transcription, a textual translation, a textual description, a visual transcription, or a visual annotation of a portion of the network resource.

16. The system of claim 15, wherein at least one visual transcription of the portion of the network resource comprises

40

a sign language translation of at least a portion of a spoken element of the network resource.

17. The system of claim 15, wherein at least one aural description of the portion of the network resource comprises a verbal description of an image or video displayed in the portion of the network resource.

18. The system of claim 15, wherein at least one aural description of the portion of the network resource comprises a verbal description of a result of an action taken on an element displayed in the portion of the network resource.

19. The system of claim 18, wherein the element comprises one of a link, button, dropdown, menu, or input control, and wherein the action taken comprises one of clicking on, typing in, or activating the element.

20. The system of claim 13, wherein the narration further comprises at least one visual effect.

21. The system of claim 20, wherein the visual effect indicates the portion of the network resource described by the second audio clip.

22. The system of claim 13, wherein determining whether a portion of the network resource has changed since the narration comprising the first audio clip was received comprises processing third-party communications associated with the portion.

23. The system of claim 13, wherein determining whether a portion of the network resource has changed since the narration comprising the first audio clip was received comprises analyzing layout tags corresponding to display areas of the network resource.

24. The system of claim 13, wherein a user of the second client computing device may continue to interact directly with the network resource during playback of at least the portion of the narration.

* * * * *

UNITED STATES PATENT AND TRADEMARK OFFICE
CERTIFICATE OF CORRECTION

PATENT NO. : 9,087,024 B1
APPLICATION NO. : 13/359028
DATED : July 21, 2015
INVENTOR(S) : Andrew Lynch Hayden

Page 1 of 1

It is certified that error appears in the above-identified patent and that said Letters Patent is hereby corrected as shown below:

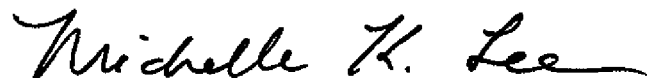
In the Specification

In column 25 at line 11, Change “above” to --above.--.

In column 34 at line 37, Change “loner” to --longer--.

In column 36 at line 30, Change “present” to --present.--.

Signed and Sealed this
Eighth Day of March, 2016

A handwritten signature in black ink, reading "Michelle K. Lee". The signature is fluid and cursive, with the first letters of each name being capitalized and prominent.

Michelle K. Lee
Director of the United States Patent and Trademark Office